

SAA-C03 Dumps

AWS Certified Solutions Architect - Associate (SAA-C03)

<https://www.certleader.com/SAA-C03-dumps.html>



NEW QUESTION 1

- (Topic 1)

A company hosts a containerized web application on a fleet of on-premises servers that process incoming requests. The number of requests is growing quickly. The on-premises servers cannot handle the increased number of requests. The company wants to move the application to AWS with minimum code changes and minimum development effort.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use AWS Fargate on Amazon Elastic Container Service (Amazon ECS) to run the containerized web application with Service Auto Scaling
- B. Use an Application Load Balancer to distribute the incoming requests.
- C. Use two Amazon EC2 instances to host the containerized web application
- D. Use an Application Load Balancer to distribute the incoming requests
- E. Use AWS Lambda with a new code that uses one of the supported languages
- F. Create multiple Lambda functions to support the load
- G. Use Amazon API Gateway as an entry point to the Lambda functions.
- H. Use a high performance computing (HPC) solution such as AWS ParallelCluster to establish an HPC cluster that can process the incoming requests at the appropriate scale.

Answer: A

Explanation:

AWS Fargate is a serverless compute engine that lets users run containers without having to manage servers or clusters of Amazon EC2 instances¹. Users can use AWS Fargate on Amazon Elastic Container Service (Amazon ECS) to run the containerized web application with Service Auto Scaling. Amazon ECS is a fully managed container orchestration service that supports both Docker and Kubernetes². Service Auto Scaling is a feature that allows users to adjust the desired number of tasks in an ECS service based on CloudWatch metrics, such as CPU utilization or request count³. Users can use AWS Fargate on Amazon ECS to migrate the application to AWS with minimum code changes and minimum development effort, as they only need to package their application in containers and specify the CPU and memory requirements.

Users can also use an Application Load Balancer to distribute the incoming requests. An Application Load Balancer is a load balancer that operates at the application layer and routes traffic to targets based on the content of the request. Users can register their ECS tasks as targets for an Application Load Balancer and configure listener rules to route requests to different target groups based on path or host headers. Users can use an Application Load Balancer to improve the availability and performance of their web application.

NEW QUESTION 2

- (Topic 1)

A company performs monthly maintenance on its AWS infrastructure. During these maintenance activities, the company needs to rotate the credentials for its Amazon ROS for MySQL databases across multiple AWS Regions

Which solution will meet these requirements with the LEAST operational overhead?

- A. Store the credentials as secrets in AWS Secrets Manager
- B. Use multi-Region secret replication for the required Regions Configure Secrets Manager to rotate the secrets on a schedule
- C. Store the credentials as secrets in AWS Systems Manager by creating a secure string parameter Use multi-Region secret replication for the required Regions Configure Systems Manager to rotate the secrets on a schedule
- D. Store the credentials in an Amazon S3 bucket that has server-side encryption (SSE) enabled Use Amazon EventBridge (Amazon CloudWatch Events) to invoke an AWS Lambda function to rotate the credentials
- E. Encrypt the credentials as secrets by using AWS Key Management Service (AWS KMS) multi-Region customer managed keys Store the secrets in an Amazon DynamoDB global table Use an AWS Lambda function to retrieve the secrets from DynamoDB Use the RDS API to rotate the secrets.

Answer: A

Explanation:

<https://aws.amazon.com/blogs/security/how-to-replicate-secrets-aws-secrets-manager-multiple-regions/>

NEW QUESTION 3

- (Topic 1)

A company is running a popular social media website. The website gives users the ability to upload images to share with other users. The company wants to make sure that the images do not contain inappropriate content. The company needs a solution that minimizes development effort.

What should a solutions architect do to meet these requirements?

- A. Use Amazon Comprehend to detect inappropriate content
- B. Use human review for low-confidence predictions.
- C. Use Amazon Rekognition to detect inappropriate content
- D. Use human review for low-confidence predictions.
- E. Use Amazon SageMaker to detect inappropriate content
- F. Use ground truth to label low-confidence predictions.
- G. Use AWS Fargate to deploy a custom machine learning model to detect inappropriate content
- H. Use ground truth to label low-confidence predictions.

Answer: B

Explanation:

<https://docs.aws.amazon.com/rekognition/latest/dg/moderation.html?pg=ln&sec=ft> <https://docs.aws.amazon.com/rekognition/latest/dg/a2i-rekognition.html>

NEW QUESTION 4

- (Topic 1)

A company runs an online marketplace web application on AWS. The application serves hundreds of thousands of users during peak hours. The company needs a scalable, near-real-time solution to share the details of millions of financial transactions with several other internal applications. Transactions also need to be processed to remove sensitive data before being stored in a document database for low-latency retrieval.

What should a solutions architect recommend to meet these requirements?

- A. Store the transactions data into Amazon DynamoDB Set up a rule in DynamoDB to remove sensitive data from every transaction upon write Use DynamoDB Streams to share the transactions data with other applications
- B. Stream the transactions data into Amazon Kinesis Data Firehose to store data in Amazon DynamoDB and Amazon S3 Use AWS Lambda integration with Kinesis Data Firehose to remove sensitive data
- C. Other applications can consume the data stored in Amazon S3
- D. Stream the transactions data into Amazon Kinesis Data Streams Use AWS Lambda integration to remove sensitive data from every transaction and then store the transactions data in Amazon DynamoDB Other applications can consume the transactions data off the Kinesis data stream.
- E. Store the batched transactions data in Amazon S3 as file
- F. Use AWS Lambda to process every file and remove sensitive data before updating the files in Amazon S3 The Lambda function then stores the data in Amazon DynamoDB Other applications can consume transaction files stored in Amazon S3.

Answer: C

Explanation:

The destination of your Kinesis Data Firehose delivery stream. Kinesis Data Firehose can send data records to various destinations, including Amazon Simple Storage Service (Amazon S3), Amazon Redshift, Amazon OpenSearch Service, and any HTTP endpoint that is owned by you or any of your third-party service providers. The following are the supported destinations:

- * Amazon OpenSearch Service
- * Amazon S3
- * Datadog
- * Dynatrace
- * Honeycomb
- * HTTP Endpoint
- * Logic Monitor
- * MongoDB Cloud
- * New Relic
- * Splunk
- * Sumo Logic <https://docs.aws.amazon.com/firehose/latest/dev/create-name.html>

<https://aws.amazon.com/kinesis/data-streams/>

Amazon Kinesis Data Streams (KDS) is a massively scalable and durable real-time data streaming service. KDS can continuously capture gigabytes of data per second from hundreds of thousands of sources such as website clickstreams, database event streams, financial transactions, social media feeds, IT logs, and location-tracking events.

NEW QUESTION 5

- (Topic 1)

A company has an application that generates a large number of files, each approximately 5 MB in size. The files are stored in Amazon S3. Company policy requires the files to be stored for 4 years before they can be deleted Immediate accessibility is always required as the files contain critical business data that is not easy to reproduce. The files are frequently accessed in the first 30 days of the object creation but are rarely accessed after the first 30 days Which storage solution is MOST cost-effective?

- A. Create an S3 bucket lifecycle policy to move Mm from S3 Standard to S3 Glacier 30 days from object creation Delete the Tiles 4 years after object creation
- B. Create an S3 bucket lifecycle policy to move tiles from S3 Standard to S3 One Zone- infrequent Access (S3 One Zone-IA] 30 days from object creatio
- C. Delete the fees 4 years after object creation
- D. Create an S3 bucket lifecycle policy to move files from S3 Standard-infrequent Access (S3 Standard -IA) 30 from object creatio
- E. Delete the ties 4 years after object creation
- F. Create an S3 bucket Lifecycle policy to move files from S3 Standard to S3 Standard- Infrequent Access (S3 Standard-IA) 30 days from object creation Move the files to S3 Glacier 4 years after object carton.

Answer: B

Explanation:

https://aws.amazon.com/s3/storage-classes/?trk=66264cd8-3b73-416c-9693-ea7cf4fe846a&sc_channel=ps&s_kwid=AL!4422!3!536452716950!p!!g!!aws%20s3%20pricing&ef_id=Cj0KcQjwnbmaBhD-ARIsAGTPcfVHUZN5_BMrzI5zBcaC8KnqpnNZvbZzqPkH6k7q4JcYO5KFLx0YYgaAm6nEALw_wcB:G:s&s_kwid=AL!4422!3!536452716950!p!!g!!aws%20s3%20pricing

NEW QUESTION 6

- (Topic 1)

A company has an application that provides marketing services to stores. The services are based on previous purchases by store customers. The stores upload transaction data to the company through SFTP, and the data is processed and analyzed to generate new marketing offers. Some of the files can exceed 200 GB in size.

Recently, the company discovered that some of the stores have uploaded files that contain personally identifiable information (PII) that should not have been included. The company wants administrators to be alerted if PII is shared again. The company also wants to automate remediation.

What should a solutions architect do to meet these requirements with the LEAST development effort?

- A. Use an Amazon S3 bucket as a secure transfer poin
- B. Use Amazon Inspector to scan me objects in the bucke
- C. If objects contain PI
- D. trigger an S3 Lifecycle policy to remove the objects that contain PII.
- E. Use an Amazon S3 bucket as a secure transfer poin
- F. Use Amazon Macie to scan the objects in the bucke
- G. If objects contain PI
- H. Use Amazon Simple Notification Service (Amazon SNS) to trigger a notification to the administrators to remove the objects mat contain PII.
- I. Implement custom scanning algorithms in an AWS Lambda functio
- J. Trigger the function when objects are loaded into the bucke
- K. It objects contain RI
- L. use Amazon Simple Notification Service (Amazon SNS) to trigger a notification to the administrators to remove the objects that contain PII.
- M. Implement custom scanning algorithms in an AWS Lambda functio
- N. Trigger the function when objects are loaded into the bucke
- O. If objects contain PI
- P. use Amazon Simple Email Service (Amazon STS) to trigger a notification to the administrators and trigger on S3 Lifecycle policy to remove the objects mot

contain PII.

Answer: B

Explanation:

To meet the requirements of detecting and alerting the administrators when PII is shared and automating remediation with the least development effort, the best approach would be to use Amazon S3 bucket as a secure transfer point and scan the objects in the bucket with Amazon Macie. Amazon Macie is a fully managed data security and data privacy service that uses machine learning and pattern matching to discover and protect sensitive data stored in Amazon S3. It can be used to classify sensitive data, monitor access to sensitive data, and automate remediation actions.

In this scenario, after uploading the files to the Amazon S3 bucket, the objects can be scanned for PII by Amazon Macie, and if it detects any PII, it can trigger an Amazon Simple Notification Service (SNS) notification to alert the administrators to remove the objects containing PII. This approach requires the least development effort, as Amazon Macie already has pre-built data classification rules that can detect PII in various formats. Hence, option B is the correct answer.

References:

? Amazon Macie User Guide: <https://docs.aws.amazon.com/macie/latest/userguide/what-is-macie.html>

? AWS Well-Architected Framework - Security Pillar: <https://docs.aws.amazon.com/wellarchitected/latest/security-pillar/welcome.html>

NEW QUESTION 7

- (Topic 1)

A company hosts more than 300 global websites and applications. The company requires a platform to analyze more than 30 TB of clickstream data each day. What should a solutions architect do to transmit and process the clickstream data?

- A. Design an AWS Data Pipeline to archive the data to an Amazon S3 bucket and run an Amazon EMR cluster with the data to generate analytics
- B. Create an Auto Scaling group of Amazon EC2 instances to process the data and send it to an Amazon S3 data lake for Amazon Redshift to use for analysis
- C. Cache the data to Amazon CloudFront: Store the data in an Amazon S3 bucket. When an object is added to the S3 bucket, run an AWS Lambda function to process the data for analysis.
- D. Collect the data from Amazon Kinesis Data Stream
- E. Use Amazon Kinesis Data Firehose to transmit the data to an Amazon S3 data lake. Load the data in Amazon Redshift for analysis

Answer: D

Explanation:

<https://aws.amazon.com/es/blogs/big-data/real-time-analytics-with-amazon-redshift-streaming-ingestion/>

NEW QUESTION 8

- (Topic 1)

A company runs a highly available image-processing application on Amazon EC2 instances in a single VPC. The EC2 instances run inside several subnets across multiple Availability Zones. The EC2 instances do not communicate with each other. However, the EC2 instances download images from Amazon S3 and upload images to Amazon S3 through a single NAT gateway. The company is concerned about data transfer charges. What is the MOST cost-effective way for the company to avoid Regional data transfer charges?

- A. Launch the NAT gateway in each Availability Zone
- B. Replace the NAT gateway with a NAT instance
- C. Deploy a gateway VPC endpoint for Amazon S3
- D. Provision an EC2 Dedicated Host to run the EC2 instances

Answer: A

Explanation:

In this scenario, the company wants to avoid regional data transfer charges while downloading and uploading images from Amazon S3. To accomplish this at the lowest cost, the NAT gateway should be launched in each availability zone that the EC2 instances are running in. This allows the EC2 instances to route traffic through the local NAT gateway instead of sending traffic across an availability zone boundary and incurring regional data transfer fees. This method will help reduce the data transfer costs since inter- Availability Zone data transfers in a single region are free of charge.

Reference:

AWS NAT Gateway documentation: <https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-gateway.html>

NEW QUESTION 9

- (Topic 1)

A company has an on-premises application that generates a large amount of time-sensitive data that is backed up to Amazon S3. The application has grown and there are user complaints about internet bandwidth limitations. A solutions architect needs to design a long-term solution that allows for both timely backups to Amazon S3 and with minimal impact on internet connectivity for internal users.

Which solution meets these requirements?

- A. Establish AWS VPN connections and proxy all traffic through a VPC gateway endpoint
- B. Establish a new AWS Direct Connect connection and direct backup traffic through this new connection.
- C. Order daily AWS Snowball devices. Load the data onto the Snowball devices and return the devices to AWS each day.
- D. Submit a support ticket through the AWS Management Console. Request the removal of S3 service limits from the account.

Answer: B

Explanation:

To address the issue of bandwidth limitations on the company's on-premises application, and to minimize the impact on internal user connectivity, a new AWS Direct Connect connection should be established to direct backup traffic through this new connection. This solution will offer a secure, high-speed connection between the company's data center and AWS, which will allow the company to transfer data quickly without consuming internet bandwidth.

Reference:

AWS Direct Connect documentation: <https://aws.amazon.com/directconnect/>

NEW QUESTION 10

- (Topic 1)

A company's website uses an Amazon EC2 instance store for its catalog of items. The company wants to make sure that the catalog is highly available and that

the catalog is stored in a durable location.

What should a solutions architect do to meet these requirements?

- A. Move the catalog to Amazon ElastiCache for Redis.
- B. Deploy a larger EC2 instance with a larger instance store.
- C. Move the catalog from the instance store to Amazon S3 Glacier Deep Archive.
- D. Move the catalog to an Amazon Elastic File System (Amazon EFS) file system.

Answer: D

Explanation:

Moving the catalog to an Amazon Elastic File System (Amazon EFS) file system provides both high availability and durability. Amazon EFS is a fully-managed, highly-available, and durable file system that is built to scale on demand. With Amazon EFS, the catalog data can be stored and accessed from multiple EC2 instances in different availability zones, ensuring high availability. Also, Amazon EFS automatically stores files redundantly within and across multiple availability zones, making it a durable storage option.

NEW QUESTION 10

- (Topic 1)

A company is implementing a new business application. The application runs on two Amazon EC2 instances and uses an Amazon S3 bucket for document storage. A solutions architect needs to ensure that the EC2 instances can access the S3 bucket.

What should the solutions architect do to meet this requirement?

- A. Create an IAM role that grants access to the S3 bucket
- B. Attach the role to the EC2 instances.
- C. Create an IAM policy that grants access to the S3 bucket
- D. Attach the policy to the EC2 instances.
- E. Create an IAM group that grants access to the S3 bucket
- F. Attach the group to the EC2 instances.
- G. Create an IAM user that grants access to the S3 bucket
- H. Attach the user account to the EC2 instances.

Answer: A

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/ec2-instance-access-s3-bucket/>

NEW QUESTION 14

- (Topic 1)

A company uses NFS to store large video files in on-premises network attached storage. Each video file ranges in size from 1MB to 500 GB. The total storage is 70 TB and is no longer growing. The company decides to migrate the video files to Amazon S3. The company must migrate the video files as soon as possible while using the least possible network bandwidth.

Which solution will meet these requirements?

- A. Create an S3 bucket Create an IAM role that has permissions to write to the S3 bucket
- B. Use the AWS CLI to copy all files locally to the S3 bucket.
- C. Create an AWS Snowball Edge job
- D. Receive a Snowball Edge device on premise
- E. Use the Snowball Edge client to transfer data to the device
- F. Return the device so that AWS can import the data into Amazon S3.
- G. Deploy an S3 File Gateway on premise
- H. Create a public service endpoint to connect to the S3 File Gateway Create an S3 bucket Create a new NFS file share on the S3 File Gateway Point the new file share to the S3 bucket
- I. Transfer the data from the existing NFS file share to the S3 File Gateway.
- J. Set up an AWS Direct Connect connection between the on-premises network and AWS
- K. Deploy an S3 File Gateway on premise
- L. Create a public virtual interface (VIF) to connect to the S3 File Gateway
- M. Create an S3 bucket
- N. Create a new NFS file share on the S3 File Gateway
- O. Point the new file share to the S3 bucket
- P. Transfer the data from the existing NFS file share to the S3 File Gateway.

Answer: B

Explanation:

The basic difference between Snowball and Snowball Edge is the capacity they provide. Snowball provides a total of 50 TB or 80 TB, out of which 42 TB or 72 TB is available, while Amazon Snowball Edge provides 100 TB, out of which 83 TB is available.

NEW QUESTION 17

- (Topic 1)

A company has a production workload that runs on 1,000 Amazon EC2 Linux instances. The workload is powered by third-party software. The company needs to patch the third-party software on all EC2 instances as quickly as possible to remediate a critical security vulnerability.

What should a solutions architect do to meet these requirements?

- A. Create an AWS Lambda function to apply the patch to all EC2 instances.
- B. Configure AWS Systems Manager Patch Manager to apply the patch to all EC2 instances.
- C. Schedule an AWS Systems Manager maintenance window to apply the patch to all EC2 instances.
- D. Use AWS Systems Manager Run Command to run a custom command that applies the patch to all EC2 instances.

Answer: B

Explanation:

<https://docs.aws.amazon.com/systems-manager/latest/userguide/about-windows-app-patching.html>

NEW QUESTION 21

- (Topic 1)

A company hosts its multi-tier applications on AWS. For compliance, governance, auditing, and security, the company must track configuration changes on its AWS resources and record a history of API calls made to these resources.

What should a solutions architect do to meet these requirements?

- A. Use AWS CloudTrail to track configuration changes and AWS Config to record API calls
- B. Use AWS Config to track configuration changes and AWS CloudTrail to record API calls
- C. Use AWS Config to track configuration changes and Amazon CloudWatch to record API calls
- D. Use AWS CloudTrail to track configuration changes and Amazon CloudWatch to record API calls

Answer: B

Explanation:

AWS Config is a fully managed service that allows the company to assess, audit, and evaluate the configurations of its AWS resources. It provides a detailed inventory of the resources in use and tracks changes to resource configurations. AWS Config can detect configuration changes and alert the company when changes occur. It also provides a historical view of changes, which is essential for compliance and governance purposes. AWS CloudTrail is a fully managed service that provides a detailed history of API calls made to the company's AWS resources. It records all API activity in the AWS account, including who made the API call, when the call was made, and what resources were affected by the call. This information is critical for security and auditing purposes, as it allows the company to investigate any suspicious activity that might occur on its AWS resources.

NEW QUESTION 24

- (Topic 1)

A company runs an on-premises application that is powered by a MySQL database. The company is migrating the application to AWS to increase the application's elasticity and availability.

The current architecture shows heavy read activity on the database during times of normal operation. Every 4 hours, the company's development team pulls a full export of the production database to populate a database in the staging environment. During this period, users experience unacceptable application latency. The development team is unable to use the staging environment until the procedure completes.

A solutions architect must recommend replacement architecture that alleviates the application latency issue. The replacement architecture also must give the development team the ability to continue using the staging environment without delay.

Which solution meets these requirements?

- A. Use Amazon Aurora MySQL with Multi-AZ Aurora Replicas for production.
- B. Populate the staging database by implementing a backup and restore process that uses the mysqldump utility.
- C. Use Amazon Aurora MySQL with Multi-AZ Aurora Replicas for production. Use database cloning to create the staging database on-demand.
- D. Use Amazon RDS for MySQL with a Multi-AZ deployment and read replicas for production. Use the standby instance for the staging database.
- E. Use Amazon RDS for MySQL with a Multi-AZ deployment and read replicas for production.
- F. Populate the staging database by implementing a backup and restore process that uses the mysqldump utility.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/aws/amazon-aurora-fast-database-cloning/>

NEW QUESTION 29

- (Topic 1)

A company uses 50 TB of data for reporting. The company wants to move this data from on-premises to AWS. A custom application in the company's data center runs a weekly data transformation job. The company plans to pause the application until the data transfer is complete and needs to begin the transfer process as soon as possible.

The data center does not have any available network bandwidth for additional workloads. A solutions architect must transfer the data and must configure the transformation job to continue to run in the AWS Cloud.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use AWS DataSync to move the data. Create a custom transformation job by using AWS Glue.
- B. Order an AWS Snowcone device to move the data. Deploy the transformation application to the device.
- C. Order an AWS Snowball Edge Storage Optimized device.
- D. Copy the data to the device.
- E. Create a custom transformation job by using AWS Glue.
- F. Order an AWS Snowball Edge Storage Optimized device that includes Amazon EC2 compute. Copy the data to the device. Create a new EC2 instance on AWS to run the transformation application.

Answer: D

Explanation:

AWS Snowball Edge is a type of Snowball device with on-board storage and compute power for select AWS capabilities. Snowball Edge can do local processing and edge-computing workloads in addition to transferring data between your local environment and the AWS Cloud¹. Users can order an AWS Snowball Edge Storage Optimized device that includes Amazon EC2 compute to move 50 TB of data from on-premises to AWS. The Storage Optimized device has 80 TB of usable storage and 40 vCPUs of compute power². Users can copy the data to the device using the AWS OpsHub graphical user interface or the Snowball client command line tool³. Users can also create and run Amazon EC2 instances on the device using Amazon Machine Images (AMIs) that are compatible with the sb1 instance type. Users can use the Snowball Edge device to transfer the data and run the transformation job locally without using any network bandwidth.

Users can also create a new EC2 instance on AWS to run the transformation application after the data transfer is complete. Amazon EC2 is a web service that provides secure, resizable compute capacity in the cloud. Users can launch an EC2 instance in the same AWS Region where they send their Snowball Edge device and choose an AMI that matches their application requirements. Users can use the EC2 instance to continue running the transformation job in the AWS Cloud.

NEW QUESTION 34

- (Topic 1)

A company provides a Voice over Internet Protocol (VoIP) service that uses UDP connections. The service consists of Amazon EC2 instances that run in an Auto Scaling group. The company has deployments across multiple AWS Regions.

The company needs to route users to the Region with the lowest latency. The company also needs automated failover between Regions.

Which solution will meet these requirements?

- A. Deploy a Network Load Balancer (NLB) and an associated target group
- B. Associate the target group with the Auto Scaling group
- C. Use the NLB as an AWS Global Accelerator endpoint in each Region.
- D. Deploy an Application Load Balancer (ALB) and an associated target group
- E. Associate the target group with the Auto Scaling group
- F. Use the ALB as an AWS Global Accelerator endpoint in each Region.
- G. Deploy a Network Load Balancer (NLB) and an associated target group
- H. Associate the target group with the Auto Scaling group
- I. Create an Amazon Route 53 latency record that points to aliases for each NL
- J. Create an Amazon CloudFront distribution that uses the latency record as an origin.
- K. Deploy an Application Load Balancer (ALB) and an associated target group
- L. Associate the target group with the Auto Scaling group
- M. Create an Amazon Route 53 weighted record that points to aliases for each AL
- N. Deploy an Amazon CloudFront distribution that uses the weighted record as an origin.

Answer: D

Explanation:

<https://aws.amazon.com/global-accelerator/faqs/>

HTTP /HTTPS - ALB ; TCP and UDP - NLB; Lowest latency routing and more throughput. Also supports failover, uses Anycast Ip addressing - Global Accelerator Caching at Edge Locations – Cloudfront

WS Global Accelerator automatically checks the health of your applications and routes user traffic only to healthy application endpoints. If the health status changes or you make configuration updates, AWS Global Accelerator reacts instantaneously to route your users to the next available endpoint..

NEW QUESTION 38

- (Topic 1)

A company's HTTP application is behind a Network Load Balancer (NLB). The NLB's target group is configured to use an Amazon EC2 Auto Scaling group with multiple EC2 instances that run the web service.

The company notices that the NLB is not detecting HTTP errors for the application. These errors require a manual restart of the EC2 instances that run the web service. The company needs to improve the application's availability without writing custom scripts or code.

What should a solutions architect do to meet these requirements?

- A. Enable HTTP health checks on the NL
- B. supplying the URL of the company's application.
- C. Add a cron job to the EC2 instances to check the local application's logs once each minute
- D. If HTTP errors are detected, the application will restart.
- E. Replace the NLB with an Application Load Balance
- F. Enable HTTP health checks by supplying the URL of the company's application
- G. Configure an Auto Scaling action to replace unhealthy instances.
- H. Create an Amazon Cloud Watch alarm that monitors the UnhealthyHostCount metric for the NL
- I. Configure an Auto Scaling action to replace unhealthy instances when the alarm is in the ALARM state.

Answer: C

Explanation:

Application availability: NLB cannot assure the availability of the application. This is because it bases its decisions solely on network and TCP-layer variables and has no awareness of the application at all. Generally, NLB determines availability based on the ability of a server to respond to ICMP ping or to correctly complete the three-way TCP handshake. ALB goes much deeper and is capable of determining availability based on not only a successful HTTP GET of a particular page but also the verification that the content is as was expected based on the input parameters.

NEW QUESTION 40

- (Topic 1)

A solutions architect is using Amazon S3 to design the storage architecture of a new digital media application. The media files must be resilient to the loss of an Availability Zone. Some files are accessed frequently while other files are rarely accessed in an unpredictable pattern. The solutions architect must minimize the costs of storing and retrieving the media files.

Which storage option meets these requirements?

- A. S3 Standard
- B. S3 Intelligent-Tiering
- C. S3 Standard-Infrequent Access (S3 Standard-IA)
- D. S3 One Zone-Infrequent Access (S3 One Zone-IA)

Answer: B

Explanation:

S3 Intelligent-Tiering - Perfect use case when you don't know the frequency of access or irregular patterns of usage.

Amazon S3 offers a range of storage classes designed for different use cases. These include S3 Standard for general-purpose storage of frequently accessed data; S3 Intelligent-Tiering for data with unknown or changing access patterns; S3 Standard- Infrequent Access (S3 Standard-IA) and S3 One Zone-Infrequent Access (S3 One Zone- IA) for long-lived, but less frequently accessed data; and Amazon S3 Glacier (S3 Glacier) and Amazon S3 Glacier Deep Archive (S3 Glacier Deep Archive) for long-term archive and digital preservation. If you have data residency requirements that can't be met by an existing AWS Region, you can use the S3 Outposts storage class to store your S3 data on-premises. Amazon S3 also offers capabilities to manage your data throughout its lifecycle. Once an S3 Lifecycle policy is set, your data will automatically transfer to a different storage class without any changes to your application.

https://aws.amazon.com/getting-started/hands-on/getting-started-using-amazon-s3-intelligent-tiering/?nc1=h_ls

NEW QUESTION 45

- (Topic 1)

A company is hosting a static website on Amazon S3 and is using Amazon Route 53 for DNS. The website is experiencing increased demand from around the world. The company must decrease latency for users who access the website.

Which solution meets these requirements MOST cost-effectively?

- A. Replicate the S3 bucket that contains the website to all AWS Region
- B. Add Route 53 geolocation routing entries.
- C. Provision accelerators in AWS Global Accelerator
- D. Associate the supplied IP addresses with the S3 bucket
- E. Edit the Route 53 entries to point to the IP addresses of the accelerators.
- F. Add an Amazon CloudFront distribution in front of the S3 bucket
- G. Edit the Route 53 entries to point to the CloudFront distribution.
- H. Enable S3 Transfer Acceleration on the bucket
- I. Edit the Route 53 entries to point to the new endpoint.

Answer: C

Explanation:

Amazon CloudFront is a content delivery network (CDN) that caches content at edge locations around the world, providing low latency and high transfer speeds to users accessing the content. Adding a CloudFront distribution in front of the S3 bucket will cache the static website's content at edge locations around the world, decreasing latency for users accessing the website. This solution is also cost-effective as it only charges for the data transfer and requests made by users accessing the content from the CloudFront edge locations. Additionally, this solution provides scalability and reliability benefits as CloudFront can automatically scale to handle increased demand and provide high availability for the website.

NEW QUESTION 49

- (Topic 1)

A company's containerized application runs on an Amazon EC2 instance. The application needs to download security certificates before it can communicate with other business applications. The company wants a highly secure solution to encrypt and decrypt the certificates in near real time. The solution also needs to store data in highly available storage after the data is encrypted.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Create AWS Secrets Manager secrets for encrypted certificate
- B. Manually update the certificates as needed
- C. Control access to the data by using fine-grained IAM access.
- D. Create an AWS Lambda function that uses the Python cryptography library to receive and perform encryption operation
- E. Store the function in an Amazon S3 bucket.
- F. Create an AWS Key Management Service (AWS KMS) customer managed key
- G. Allow the EC2 role to use the KMS key for encryption operation
- H. Store the encrypted data on Amazon S3.
- I. Create an AWS Key Management Service (AWS KMS) customer managed key
- J. Allow the EC2 role to use the KMS key for encryption operation
- K. Store the encrypted data on Amazon Elastic Block Store (Amazon EBS) volumes.

Answer: D

NEW QUESTION 52

- (Topic 1)

A company runs an ecommerce application on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. The Auto Scaling group scales based on CPU utilization metrics. The ecommerce application stores the transaction data in a MySQL 8.0 database that is hosted on a large EC2 instance.

The database's performance degrades quickly as application load increases. The application handles more read requests than write transactions. The company wants a solution that will automatically scale the database to meet the demand of unpredictable read workloads while maintaining high availability.

Which solution will meet these requirements?

- A. Use Amazon Redshift with a single node for leader and compute functionality.
- B. Use Amazon RDS with a Single-AZ deployment. Configure Amazon RDS to add reader instances in a different Availability Zone.
- C. Use Amazon Aurora with a Multi-AZ deployment
- D. Configure Aurora Auto Scaling with Aurora Replicas.
- E. Use Amazon ElastiCache for Memcached with EC2 Spot Instances.

Answer: C

Explanation:

AURORA is 5x performance improvement over MySQL on RDS and handles more read requests than write,; maintaining high availability = Multi-AZ deployment

NEW QUESTION 56

- (Topic 1)

A company has a website hosted on AWS. The website is behind an Application Load Balancer (ALB) that is configured to handle HTTP and HTTPS separately. The company wants to forward all requests to the website so that the requests will use HTTPS.

What should a solutions architect do to meet this requirement?

- A. Update the ALB's network ACL to accept only HTTPS traffic
- B. Create a rule that replaces the HTTP in the URL with HTTPS.
- C. Create a listener rule on the ALB to redirect HTTP traffic to HTTPS.
- D. Replace the ALB with a Network Load Balancer configured to use Server Name Indication (SNI).

Answer: C

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/elb-redirect-http-to-https-using-alb/>

How can I redirect HTTP requests to HTTPS using an Application Load Balancer? Last updated: 2020-10-30 I want to redirect HTTP requests to HTTPS using Application Load Balancer listener rules. How can I do this? Resolution Reference: <https://aws.amazon.com/premiumsupport/knowledge-center/elb-redirect-http-to-https-using-alb/>

NEW QUESTION 61

- (Topic 1)

A company is hosting a web application on AWS using a single Amazon EC2 instance that stores user-uploaded documents in an Amazon EBS volume. For better scalability and availability, the company duplicated the architecture and created a second EC2 instance and EBS volume in another Availability Zone placing both behind an Application Load Balancer. After completing this change, users reported that, each time they refreshed the website, they could see one subset of their documents or the other, but never all of the documents at the same time.

What should a solutions architect propose to ensure users see all of their documents at once?

- A. Copy the data so both EBS volumes contain all the documents.
- B. Configure the Application Load Balancer to direct a user to the server with the documents
- C. Copy the data from both EBS volumes to Amazon EFS. Modify the application to save new documents to Amazon EFS
- D. Configure the Application Load Balancer to send the request to both servers. Return each document from the correct server.

Answer: C

Explanation:

<https://docs.aws.amazon.com/efs/latest/ug/how-it-works.html#how-it-works-ec2>

NEW QUESTION 64

- (Topic 1)

An image-processing company has a web application that users use to upload images. The application uploads the images into an Amazon S3 bucket. The company has set up S3 event notifications to publish the object creation events to an Amazon Simple Queue Service (Amazon SQS) standard queue. The SQS queue serves as the event source for an AWS Lambda function that processes the images and sends the results to users through email.

Users report that they are receiving multiple email messages for every uploaded image. A solutions architect determines that SQS messages are invoking the Lambda function more than once, resulting in multiple email messages.

What should the solutions architect do to resolve this issue with the LEAST operational overhead?

- A. Set up long polling in the SQS queue by increasing the ReceiveMessage wait time to 30 seconds.
- B. Change the SQS standard queue to an SQS FIFO queue
- C. Use the message deduplication ID to discard duplicate messages.
- D. Increase the visibility timeout in the SQS queue to a value that is greater than the total of the function timeout and the batch window timeout.
- E. Modify the Lambda function to delete each message from the SQS queue immediately after the message is read before processing.

Answer: C

NEW QUESTION 65

- (Topic 1)

A company has a data ingestion workflow that consists of the following:

? An Amazon Simple Notification Service (Amazon SNS) topic for notifications about new data deliveries

? An AWS Lambda function to process the data and record metadata

The company observes that the ingestion workflow fails occasionally because of network connectivity issues. When such a failure occurs, the Lambda function does not ingest the corresponding data unless the company manually reruns the job.

Which combination of actions should a solutions architect take to ensure that the Lambda function ingests all data in the future? (Select TWO.)

- A. Configure the Lambda function in multiple Availability Zones.
- B. Create an Amazon Simple Queue Service (Amazon SQS) queue, and subscribe it to the SNS topic.
- C. Increase the CPU and memory that are allocated to the Lambda function.
- D. Increase provisioned throughput for the Lambda function.
- E. Modify the Lambda function to read from an Amazon Simple Queue Service (Amazon SQS) queue

Answer: BE

Explanation:

To ensure that the Lambda function ingests all data in the future despite occasional network connectivity issues, the following actions should be taken:

? Create an Amazon Simple Queue Service (SQS) queue and subscribe it to the SNS topic. This allows for decoupling of the notification and processing, so that even if the processing Lambda function fails, the message remains in the queue for further processing later.

? Modify the Lambda function to read from the SQS queue instead of directly from SNS. This decoupling allows for retries and fault tolerance and ensures that all messages are processed by the Lambda function.

Reference:

AWS SNS documentation: <https://aws.amazon.com/sns/> AWS SQS documentation: <https://aws.amazon.com/sqs/>

AWS Lambda documentation: <https://aws.amazon.com/lambda/>

NEW QUESTION 66

- (Topic 1)

A company has applications that run on Amazon EC2 instances in a VPC. One of the applications needs to call the Amazon S3 API to store and read objects. According to the company's security regulations, no traffic from the applications is allowed to travel across the internet.

Which solution will meet these requirements?

- A. Configure an S3 interface endpoint.
- B. Configure an S3 gateway endpoint.
- C. Create an S3 bucket in a private subnet.
- D. Create an S3 bucket in the same Region as the EC2 instance.

Answer: B

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/privatelink-interface-endpoints.html#types-of-vpc-endpoints-for-s3>
<https://docs.aws.amazon.com/vpc/latest/userguide/vpc-endpoints-s3.html>

NEW QUESTION 69

- (Topic 1)

A company recently launched Linux-based application instances on Amazon EC2 in a private subnet and launched a Linux-based bastion host on an Amazon EC2 instance in a public subnet of a VPC. A solutions architect needs to connect from the on-premises network, through the company's internet connection to the bastion host and to the application servers. The solutions architect must make sure that the security groups of all the EC2 instances will allow that access.

Which combination of steps should the solutions architect take to meet these requirements? (Select TWO)

- A. Replace the current security group of the bastion host with one that only allows inbound access from the application instances.
- B. Replace the current security group of the bastion host with one that only allows inbound access from the internal IP range for the company.
- C. Replace the current security group of the bastion host with one that only allows inbound access from the external IP range for the company.
- D. Replace the current security group of the application instances with one that allows inbound SSH access from only the private IP address of the bastion host.
- E. Replace the current security group of the application instances with one that allows inbound SSH access from only the public IP address of the bastion host.

Answer: CD

Explanation:

<https://digitalcloud.training/ssh-into-ec2-in-private-subnet/>

NEW QUESTION 70

- (Topic 1)

A company is running an SMB file server in its data center. The file server stores large files that are accessed frequently for the first few days after the files are created. After 7 days the files are rarely accessed.

The total data size is increasing and is close to the company's total storage capacity. A solutions architect must increase the company's available storage space without losing low-latency access to the most recently accessed files. The solutions architect must also provide file lifecycle management to avoid future storage issues.

Which solution will meet these requirements?

- A. Use AWS DataSync to copy data that is older than 7 days from the SMB file server to AWS.
- B. Create an Amazon S3 File Gateway to extend the company's storage space.
- C. Create an S3 Lifecycle policy to transition the data to S3 Glacier Deep Archive after 7 days.
- D. Create an Amazon FSx for Windows File Server file system to extend the company's storage space.
- E. Install a utility on each user's computer to access Amazon S3. Create an S3 Lifecycle policy to transition the data to S3 Glacier Flexible Retrieval after 7 days.

Answer: B

Explanation:

Amazon S3 File Gateway is a hybrid cloud storage service that enables on-premises applications to seamlessly use Amazon S3 cloud storage. It provides a file interface to Amazon S3 and supports SMB and NFS protocols. It also supports S3 Lifecycle policies that can automatically transition data from S3 Standard to S3 Glacier Deep Archive after a specified period of time. This solution will meet the requirements of increasing the company's available storage space without losing low-latency access to the most recently accessed files and providing file lifecycle management to avoid future storage issues.

Reference:

<https://docs.aws.amazon.com/storagegateway/latest/userguide/WhatIsStorageGateway.html>

NEW QUESTION 71

- (Topic 1)

A company hosts an application on multiple Amazon EC2 instances. The application processes messages from an Amazon SQS queue, writes to an Amazon RDS table, and deletes the message from the queue. Occasional duplicate records are found in the RDS table. The SQS queue does not contain any duplicate messages.

What should a solutions architect do to ensure messages are being processed once only?

- A. Use the CreateQueue API call to create a new queue.
- B. Use the AddPermission API call to add appropriate permissions.
- C. Use the ReceiveMessage API call to set an appropriate wait time.
- D. Use the ChangeMessageVisibility API call to increase the visibility timeout.

Answer: D

Explanation:

The visibility timeout begins when Amazon SQS returns a message. During this time, the consumer processes and deletes the message. However, if the consumer fails before deleting the message and your system doesn't call the DeleteMessage action for that message before the visibility timeout expires, the message becomes visible to other consumers and the message is received again. If a message must be received only once, your consumer should delete it within the duration of the visibility timeout. <https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-visibility-timeout.html>

Keyword: SQS queue writes to an Amazon RDS. From this, Option D is the best suite & other options ruled out [Option A - You can't introduce one more Queue in the existing one; Option B - only Permission & Option C - Only Retrieves Messages]. FIFO queues are designed to never introduce duplicate messages. However, your message producer might introduce duplicates in certain scenarios: for example, if the producer sends a message, does not receive a response, and then resends the same message. Amazon SQS APIs provide deduplication functionality that prevents your message producer from sending duplicates. Any duplicates introduced by the message producer are removed within a 5-minute deduplication interval. For standard queues, you might occasionally receive a duplicate copy of a message (at-least- once delivery). If you use a standard queue, you must design your applications to be idempotent (that is, they must not be affected adversely when processing the same message more than once).

NEW QUESTION 75

- (Topic 1)

A solutions architect is designing the cloud architecture for a new application being deployed on AWS. The process should run in parallel while adding and removing application nodes as needed based on the number of jobs to be processed. The processor application is stateless. The solutions architect must ensure that the application is loosely coupled and the job items are durably stored.

Which design should the solutions architect use?

- A. Create an Amazon SNS topic to send the jobs that need to be processed Create an Amazon Machine Image (AMI) that consists of the processor application Create a launch configuration that uses the AMI Create an Auto Scaling group using the launch configuration Set the scaling policy for the Auto Scaling group to add and remove nodes based on CPU usage
- B. Create an Amazon SQS queue to hold the jobs that need to be processed Create an Amazon Machine image (AMI) that consists of the processor application Create a launch configuration that uses the AMI Create an Auto Scaling group using the launch configuration Set the scaling policy for the Auto Scaling group to add and remove nodes based on network usage
- C. Create an Amazon SQS queue to hold the jobs that need to be processed Create an Amazon Machine image (AMI) that consists of the processor application Create a launch template that uses the AMI Create an Auto Scaling group using the launch template Set the scaling policy for the Auto Scaling group to add and remove nodes based on the number of items in the SQS queue
- D. Create an Amazon SNS topic to send the jobs that need to be processed Create an Amazon Machine Image (AMI) that consists of the processor application Create a launch template that uses the AMI Create an Auto Scaling group using the launch template Set the scaling policy for the Auto Scaling group to add and remove nodes based on the number of messages published to the SNS topic

Answer: C

Explanation:

"Create an Amazon SQS queue to hold the jobs that need to be processed. Create an Amazon EC2 Auto Scaling group for the compute application. Set the scaling policy for the Auto Scaling group to add and remove nodes based on the number of items in the SQS queue"

In this case we need to find a durable and loosely coupled solution for storing jobs. Amazon SQS is ideal for this use case and can be configured to use dynamic scaling based on the number of jobs waiting in the queue. To configure this scaling you can use the backlog per instance metric with the target value being the acceptable backlog per instance to maintain. You can calculate these numbers as follows: Backlog per instance: To calculate your backlog per instance, start with the ApproximateNumberOfMessages queue attribute to determine the length of the SQS queue

NEW QUESTION 79

- (Topic 1)

A company needs to review its AWS Cloud deployment to ensure that its Amazon S3 buckets do not have unauthorized configuration changes.

What should a solutions architect do to accomplish this goal?

- A. Turn on AWS Config with the appropriate rules.
- B. Turn on AWS Trusted Advisor with the appropriate checks.
- C. Turn on Amazon Inspector with the appropriate assessment template.
- D. Turn on Amazon S3 server access logging
- E. Configure Amazon EventBridge (Amazon Cloud Watch Events).

Answer: A

Explanation:

To ensure that Amazon S3 buckets do not have unauthorized configuration changes, a solutions architect should turn on AWS Config with the appropriate rules. AWS Config is a service that allows users to audit and assess their AWS resource configurations for compliance with industry standards and internal policies. It provides a detailed view of the resources and their configurations, including information on how the resources are related to each other. By turning on AWS Config with the appropriate rules, users can identify and remediate unauthorized configuration changes to their Amazon S3 buckets.

NEW QUESTION 82

- (Topic 1)

A company collects temperature, humidity, and atmospheric pressure data in cities across multiple continents. The average volume of data collected per site each day is 500 GB. Each site has a high-speed internet connection. The company's weather forecasting applications are based in a single Region and analyze the data daily.

What is the FASTEST way to aggregate data from all of these global sites?

- A. Enable Amazon S3 Transfer Acceleration on the destination bucket
- B. Use multipart uploads to directly upload site data to the destination bucket.
- C. Upload site data to an Amazon S3 bucket in the closest AWS Region
- D. Use S3 cross-Region replication to copy objects to the destination bucket.
- E. Schedule AWS Snowball jobs daily to transfer data to the closest AWS Region
- F. Use S3 cross-Region replication to copy objects to the destination bucket.
- G. Upload the data to an Amazon EC2 instance in the closest Region
- H. Store the data in an Amazon Elastic Block Store (Amazon EBS) volume
- I. Once a day take an EBS snapshot and copy it to the centralized Region
- J. Restore the EBS volume in the centralized Region and run an analysis on the data daily.

Answer: A

Explanation:

You might want to use Transfer Acceleration on a bucket for various reasons, including the following:

You have customers that upload to a centralized bucket from all over the world. You transfer gigabytes to terabytes of data on a regular basis across continents.

You are unable to utilize all of your available bandwidth over the Internet when uploading to Amazon S3.

<https://docs.aws.amazon.com/AmazonS3/latest/dev/transfer-acceleration.html>

[https://aws.amazon.com/s3/transfer-acceleration/#:~:text=S3%20Transfer%20Acceleration%20\(S3TA\)%20reduces,to%20S3%20for%20remote%20applications:](https://aws.amazon.com/s3/transfer-acceleration/#:~:text=S3%20Transfer%20Acceleration%20(S3TA)%20reduces,to%20S3%20for%20remote%20applications:)

"Amazon S3 Transfer Acceleration can speed up content transfers to and from Amazon S3 by as much as 50-500% for long-distance transfer of larger objects.

Customers who have either web or mobile applications with widespread users or applications hosted far away from their S3 bucket can experience long and variable upload and download speeds over the Internet"

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/mpuoverview.html>

"Improved throughput - You can upload parts in parallel to improve throughput."

NEW QUESTION 86

- (Topic 1)

A company is building an application in the AWS Cloud. The application will store data in Amazon S3 buckets in two AWS Regions. The company must use an AWS Key Management Service (AWS KMS) customer managed key to encrypt all data that is stored in the S3 buckets. The data in both S3 buckets must be encrypted and decrypted with the same KMS key. The data and the key must be stored in each of the two Regions.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Create an S3 bucket in each Region Configure the S3 buckets to use server-side encryption with Amazon S3 managed encryption keys (SSE-S3) Configure replication between the S3 buckets.
- B. Create a customer managed multi-Region KMS key
- C. Create an S3 bucket in each Region
- D. Configure replication between the S3 buckets
- E. Configure the application to use the KMS key with client-side encryption.
- F. Create a customer managed KMS key and an S3 bucket in each Region Configure the S3 buckets to use server-side encryption with Amazon S3 managed encryption keys (SSE-S3) Configure replication between the S3 buckets.
- G. Create a customer managed KMS key and an S3 bucket in each Region Configure the S3 buckets to use server-side encryption with AWS KMS keys (SSE-KMS) Configure replication between the S3 buckets.

Answer: B

Explanation:

From <https://docs.aws.amazon.com/kms/latest/developerguide/custom-key-store-overview.html>

For most users, the default AWS KMS key store, which is protected by FIPS 140-2 validated cryptographic modules, fulfills their security requirements. There is no need to add an extra layer of maintenance responsibility or a dependency on an additional service. However, you might consider creating a custom key store if your organization has any of the following requirements: Key material cannot be stored in a shared environment. Key material must be subject to a secondary, independent audit path. The HSMs that generate and store key material must be certified at FIPS 140-2 Level 3.

<https://docs.aws.amazon.com/kms/latest/developerguide/custom-key-store-overview.html>

<https://docs.aws.amazon.com/kms/latest/developerguide/multi-region-keys-overview.html>

NEW QUESTION 88

- (Topic 1)

An Amazon EC2 administrator created the following policy associated with an IAM group containing several users

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "ec2:TerminateInstances",
      "Resource": "*",
      "Condition": {
        "IpAddress": {
          "aws:SourceIp": "10.100.100.0/24"
        }
      }
    },
    {
      "Effect": "Deny",
      "Action": "ec2:*",
      "Resource": "*",
      "Condition": {
        "StringNotEquals": {
          "ec2:Region": "us-east-1"
        }
      }
    }
  ]
}
```

What is the effect of this policy?

- A. Users can terminate an EC2 instance in any AWS Region except us-east-1.
- B. Users can terminate an EC2 instance with the IP address 10.100.100.1 in the us-east-1 Region
- C. Users can terminate an EC2 instance in the us-east-1 Region when the user's source IP is 10.100.100.254.
- D. Users cannot terminate an EC2 instance in the us-east-1 Region when the user's source IP is 10.100.100.254

Answer: C

Explanation:

As the policy prevents anyone from doing any EC2 action on any region except us-east-1 and allows only users with source IP 10.100.100.0/24 to terminate instances. So user with source IP 10.100.100.254 can terminate instances in us-east-1 region.

NEW QUESTION 93

- (Topic 1)

A company is planning to use an Amazon DynamoDB table for data storage. The company is concerned about cost optimization. The table will not be used on most mornings. In the evenings, the read and write traffic will often be unpredictable. When traffic spikes occur, they will happen very quickly. What should a solutions architect recommend?

- A. Create a DynamoDB table in on-demand capacity mode.
- B. Create a DynamoDB table with a global secondary index.
- C. Create a DynamoDB table with provisioned capacity and auto scaling.
- D. Create a DynamoDB table in provisioned capacity mode, and configure it as a global table.

Answer: A

NEW QUESTION 95

- (Topic 1)

A solutions architect is designing a new hybrid architecture to extend a company's on-premises infrastructure to AWS. The company requires a highly available connection with consistent low latency to an AWS Region. The company needs to minimize costs and is willing to accept slower traffic if the primary connection fails.

What should the solutions architect do to meet these requirements?

- A. Provision an AWS Direct Connect connection to a Region. Provision a VPN connection as a backup if the primary Direct Connect connection fails.
- B. Provision a VPN tunnel connection to a Region for private connectivity.
- C. Provision a second VPN tunnel for private connectivity and as a backup if the primary VPN connection fails.
- D. Provision an AWS Direct Connect connection to a Region. Provision a second Direct Connect connection to the same Region as a backup if the primary Direct Connect connection fails.
- E. Provision an AWS Direct Connect connection to a Region. Use the Direct Connect failover attribute from the AWS CLI to automatically create a backup connection if the primary Direct Connect connection fails.

Answer: A

Explanation:

"In some cases, this connection alone is not enough. It is always better to guarantee a fallback connection as the backup of DX. There are several options, but implementing it with an AWS Site-To-Site VPN is a real cost-effective solution that can be exploited to reduce costs or, in the meantime, wait for the setup of a second DX." <https://www.proud2becloud.com/hybrid-cloud-networking-backup-aws-direct-connect-network-connection-with-aws-site-to-site-vpn/>

NEW QUESTION 100

- (Topic 1)

A company needs to configure a real-time data ingestion architecture for its application. The company needs an API, a process that transforms data as the data is streamed, and a storage solution for the data.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Deploy an Amazon EC2 instance to host an API that sends data to an Amazon Kinesis data stream.
- B. Create an Amazon Kinesis Data Firehose delivery stream that uses the Kinesis data stream as a data source.
- C. Use AWS Lambda functions to transform the data.
- D. Use the Kinesis Data Firehose delivery stream to send the data to Amazon S3.
- E. Deploy an Amazon EC2 instance to host an API that sends data to AWS Glue.
- F. Stop source/destination checking on the EC2 instance.
- G. Use AWS Glue to transform the data and to send the data to Amazon S3.
- H. Configure an Amazon API Gateway API to send data to an Amazon Kinesis data stream.
- I. Create an Amazon Kinesis Data Firehose delivery stream that uses the Kinesis data stream as a data source.
- J. Use AWS Lambda functions to transform the data.
- K. Use the Kinesis Data Firehose delivery stream to send the data to Amazon S3.
- L. Configure an Amazon API Gateway API to send data to AWS Glue.
- M. Use AWS Lambda functions to transform the data.
- N. Use AWS Glue to send the data to Amazon S3.

Answer: C

NEW QUESTION 101

- (Topic 1)

A company has an Amazon S3 bucket that contains critical data. The company must protect the data from accidental deletion.

Which combination of steps should a solutions architect take to meet these requirements?

(Choose two.)

- A. Enable versioning on the S3 bucket.
- B. Enable MFA Delete on the S3 bucket.
- C. Create a bucket policy on the S3 bucket.
- D. Enable default encryption on the S3 bucket.
- E. Create a lifecycle policy for the objects in the S3 bucket.

Answer: AB

Explanation:

To protect data in an S3 bucket from accidental deletion, versioning should be enabled, which enables you to preserve, retrieve, and restore every version of every object in an S3 bucket. Additionally, enabling MFA (multi-factor authentication) Delete on the S3 bucket adds an extra layer of protection by requiring an authentication token in addition to the user's access keys to delete objects in the bucket.

Reference:

AWS S3 Versioning documentation: <https://docs.aws.amazon.com/AmazonS3/latest/dev/Versioning.html>

AWS S3 MFA Delete documentation: <https://docs.aws.amazon.com/AmazonS3/latest/dev/UsingMFADelete.html>

NEW QUESTION 103

- (Topic 1)

A company has more than 5 TB of file data on Windows file servers that run on premises. Users and applications interact with the data each day. The company is moving its Windows workloads to AWS. As the company continues this process, the company requires access to AWS and on-premises file storage with minimum latency. The company needs a solution that minimizes operational overhead and requires no significant changes to the existing file access patterns. The company uses an AWS Site-to-Site VPN connection for connectivity to AWS. What should a solutions architect do to meet these requirements?

- A. Deploy and configure Amazon FSx for Windows File Server on AWS.
- B. Move the on-premises file data to FSx for Windows File Server.
- C. Reconfigure the workloads to use FSx for Windows File Server on AWS.
- D. Deploy and configure an Amazon S3 File Gateway on premises. Move the on-premises file data to the S3 File Gateway. Reconfigure the on-premises workloads and the cloud workloads to use the S3 File Gateway.
- E. Deploy and configure an Amazon S3 File Gateway on premises. Move the on-premises file data to Amazon S3. Reconfigure the workloads to use either Amazon S3 directly or the S3 File Gateway, depending on each workload's location.
- F. Deploy and configure Amazon FSx for Windows File Server on AWS. Deploy and configure an Amazon FSx File Gateway on premises. Move the on-premises file data to the FSx File Gateway. Configure the cloud workloads to use FSx for Windows File Server on AWS. Configure the on-premises workloads to use the FSx File Gateway.

Answer: D

Explanation:

<https://docs.aws.amazon.com/filegateway/latest/filefsxw/what-is-file-fsxw.html>

To meet the requirements of the company to have access to both AWS and on-premises file storage with minimum latency, a hybrid cloud architecture can be used. One solution is to deploy and configure Amazon FSx for Windows File Server on AWS, which provides fully managed Windows file servers. The on-premises file data can be moved to the FSx File Gateway, which can act as a bridge between on-premises and AWS file storage. The cloud workloads can be configured to use FSx for Windows File Server on AWS, while the on-premises workloads can be configured to use the FSx File Gateway. This solution minimizes operational overhead and requires no significant changes to the existing file access patterns. The connectivity between on-premises and AWS can be established using an AWS Site-to-Site VPN connection.

Reference:

AWS FSx for Windows File Server: <https://aws.amazon.com/fsx/windows/> AWS FSx File Gateway: <https://aws.amazon.com/fsx/file-gateway/>

AWS Site-to-Site VPN: <https://aws.amazon.com/vpn/site-to-site-vpn/>

NEW QUESTION 108

- (Topic 1)

A company's dynamic website is hosted using on-premises servers in the United States. The company is launching its product in Europe, and it wants to optimize site loading times for new European users. The site's backend must remain in the United States. The product is being launched in a few days, and an immediate solution is needed.

What should the solutions architect recommend?

- A. Launch an Amazon EC2 instance in us-east-1 and migrate the site to it.
- B. Move the website to Amazon S3. Use cross-Region replication between Regions.
- C. Use Amazon CloudFront with a custom origin pointing to the on-premises servers.
- D. Use an Amazon Route 53 geo-proximity routing policy pointing to on-premises servers.

Answer: C

Explanation:

<https://aws.amazon.com/pt/blogs/aws/amazon-cloudfront-support-for-custom-origins/>

You can now create a CloudFront distribution using a custom origin. Each distribution can point to an S3 or to a custom origin. This could be another storage service, or it could be something more interesting and more dynamic, such as an EC2 instance or even an Elastic Load Balancer.

NEW QUESTION 111

- (Topic 1)

A company that hosts its web application on AWS wants to ensure all Amazon EC2 instances, Amazon RDS DB instances, and Amazon Redshift clusters are configured with tags. The company wants to minimize the effort of configuring and operating this check.

What should a solutions architect do to accomplish this?

- A. Use AWS Config rules to define and detect resources that are not properly tagged.
- B. Use Cost Explorer to display resources that are not properly tagged.
- C. Tag those resources manually.
- D. Write API calls to check all resources for proper tag allocation.
- E. Periodically run the code on an EC2 instance.
- F. Write API calls to check all resources for proper tag allocation.
- G. Schedule an AWS Lambda function through Amazon CloudWatch to periodically run the code.

Answer: A

Explanation:

To ensure all Amazon EC2 instances, Amazon RDS DB instances, and Amazon Redshift clusters are configured with tags, a solutions architect should use AWS Config rules to define and detect resources that are not properly tagged. AWS Config rules are a set of customizable rules that AWS Config uses to evaluate AWS resource configurations for compliance with best practices and company policies. Using AWS Config rules can minimize the effort of configuring and operating this check because it automates the process of identifying non-compliant resources and notifying the responsible teams. Reference:

AWS Config Developer Guide: AWS Config Rules (https://docs.aws.amazon.com/config/latest/developerguide/evaluate-config_use-managed_rules.html)

NEW QUESTION 113

- (Topic 1)

A company is developing a two-tier web application on AWS. The company's developers have deployed the application on an Amazon EC2 instance that connects directly to a backend Amazon RDS database. The company must not hardcode database credentials in the application. The company must also implement a solution to automatically rotate the database credentials on a regular basis.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Store the database credentials in the instance metadata
- B. Use Amazon EventBridge (Amazon CloudWatch Events) rules to run a scheduled AWS Lambda function that updates the RDS credentials and instance metadata at the same time.
- C. Store the database credentials in a configuration file in an encrypted Amazon S3 bucket
- D. Use Amazon EventBridge (Amazon CloudWatch Events) rules to run a scheduled AWS Lambda function that updates the RDS credentials and the credentials in the configuration file at the same time
- E. Use S3 Versioning to ensure the ability to fall back to previous values.
- F. Store the database credentials as a secret in AWS Secrets Manager
- G. Turn on automatic rotation for the secret
- H. Attach the required permission to the EC2 role to grant access to the secret.
- I. Store the database credentials as encrypted parameters in AWS Systems Manager Parameter Store
- J. Turn on automatic rotation for the encrypted parameter
- K. Attach the required permission to the EC2 role to grant access to the encrypted parameters.

Answer: C

Explanation:

https://docs.aws.amazon.com/secretsmanager/latest/userguide/create_database_secret.html

NEW QUESTION 115

- (Topic 1)

A company is deploying a new public web application to AWS. The application will run behind an Application Load Balancer (ALB). The application needs to be encrypted at the edge with an SSL/TLS certificate that is issued by an external certificate authority (CA). The certificate must be rotated each year before the certificate expires.

What should a solutions architect do to meet these requirements?

- A. Use AWS Certificate Manager (ACM) to issue an SSL/TLS certificate
- B. Apply the certificate to the ALB
- C. Use the managed renewal feature to automatically rotate the certificate.
- D. Use AWS Certificate Manager (ACM) to issue an SSL/TLS certificate
- E. Import the key material from the certificate
- F. Apply the certificate to the ALB
- G. Use the managed renewal feature to automatically rotate the certificate.
- H. Use AWS Certificate Manager (ACM) Private Certificate Authority to issue an SSL/TLS certificate from the root CA
- I. Apply the certificate to the ALB
- J. Use the managed renewal feature to automatically rotate the certificate.
- K. Use AWS Certificate Manager (ACM) to import an SSL/TLS certificate
- L. Apply the certificate to the ALB
- M. Use Amazon EventBridge (Amazon CloudWatch Events) to send a notification when the certificate is nearing expiration
- N. Rotate the certificate manually.

Answer: D

Explanation:

https://www.amazonaws.cn/en/certificate-manager/faqs/#Managed_renewal_and_deployment

NEW QUESTION 118

- (Topic 1)

A company is using a SQL database to store movie data that is publicly accessible. The database runs on an Amazon RDS Single-AZ DB instance. A script runs queries at random intervals each day to record the number of new movies that have been added to the database. The script must report a final total during business hours. The company's development team notices that the database performance is inadequate for development tasks when the script is running. A solutions architect must recommend a solution to resolve this issue. Which solution will meet this requirement with the LEAST operational overhead?

- A. Modify the DB instance to be a Multi-AZ deployment
- B. Create a read replica of the database. Configure the script to query only the read replica.
- C. Instruct the development team to manually export the entries in the database at the end of each day.
- D. Use Amazon ElastiCache to cache the common queries that the script runs against the database.

Answer: B

NEW QUESTION 123

- (Topic 1)

A survey company has gathered data for several years from across the United States. The company hosts the data in an Amazon S3 bucket that is 3 TB in size and growing. The company has started to share the data with a European marketing firm that has S3 buckets. The company wants to ensure that its data transfer costs remain as low as possible.

Which solution will meet these requirements?

- A. Configure the Requester Pays feature on the company's S3 bucket
- B. Configure S3 Cross-Region Replication from the company's S3 bucket to one of the marketing firm's S3 buckets.
- C. Configure cross-account access for the marketing firm so that the marketing firm has access to the company's S3 bucket.
- D. Configure the company's S3 bucket to use S3 Intelligent-Tiering. Sync the S3 bucket to one of the marketing firm's S3 buckets.

Answer: A

Explanation:

"Typically, you configure buckets to be Requester Pays buckets when you want to share data but not incur charges associated with others accessing the data. For example, you might use Requester Pays buckets when making available large datasets, such as zip code directories, reference data, geospatial information, or web crawling data." <https://docs.aws.amazon.com/AmazonS3/latest/userguide/RequesterPaysBuckets.html>

NEW QUESTION 128

- (Topic 1)

A company is migrating applications to AWS. The applications are deployed in different accounts. The company manages the accounts centrally by using AWS Organizations. The company's security team needs a single sign-on (SSO) solution across all the company's accounts. The company must continue managing the users and groups in its on-premises self-managed Microsoft Active Directory.

Which solution will meet these requirements?

- A. Enable AWS Single Sign-On (AWS SSO) from the AWS SSO console
- B. Create a one-way forest trust or a one-way domain trust to connect the company's self-managed Microsoft Active Directory with AWS SSO by using AWS Directory Service for Microsoft Active Directory.
- C. Enable AWS Single Sign-On (AWS SSO) from the AWS SSO console
- D. Create a two-way forest trust to connect the company's self-managed Microsoft Active Directory with AWS SSO by using AWS Directory Service for Microsoft Active Directory.
- E. Use AWS Directory Service
- F. Create a two-way trust relationship with the company's self-managed Microsoft Active Directory.
- G. Deploy an identity provider (IdP) on premise
- H. Enable AWS Single Sign-On (AWS SSO) from the AWS SSO console.

Answer: A

Explanation:

To provide single sign-on (SSO) across all the company's accounts while continuing to manage users and groups in its on-premises self-managed Microsoft Active Directory, the solution is to enable AWS Single Sign-On (SSO) from the AWS SSO console and create a one-way forest trust or a one-way domain trust to connect the company's self-managed Microsoft Active Directory with AWS SSO by using AWS Directory Service for Microsoft Active Directory. This solution is described in the AWS documentation

NEW QUESTION 131

- (Topic 1)

A company is designing an application where users upload small files into Amazon S3. After a user uploads a file, the file requires one-time simple processing to transform the data and save the data in JSON format for later analysis.

Each file must be processed as quickly as possible after it is uploaded. Demand will vary. On some days, users will upload a high number of files. On other days, users will upload a few files or no files.

Which solution meets these requirements with the LEAST operational overhead?

- A. Configure Amazon EMR to read text files from Amazon S3. Run processing scripts to transform the data
- B. Store the resulting JSON file in an Amazon Aurora DB cluster.
- C. Configure Amazon S3 to send an event notification to an Amazon Simple Queue Service (Amazon SQS) queue
- D. Use Amazon EC2 instances to read from the queue and process the data
- E. Store the resulting JSON file in Amazon DynamoDB.
- F. Configure Amazon S3 to send an event notification to an Amazon Simple Queue Service (Amazon SQS) queue
- G. Use an AWS Lambda function to read from the queue and process the data
- H. Store the resulting JSON file in Amazon DynamoDB
- I. Most Voted
- J. Configure Amazon EventBridge (Amazon CloudWatch Events) to send an event to Amazon Kinesis Data Streams when a new file is uploaded
- K. Use an AWS Lambda function to consume the event from the stream and process the data
- L. Store the resulting JSON file in Amazon Aurora DB cluster.

Answer: C

Explanation:

Amazon S3 sends event notifications about S3 buckets (for example, object created, object removed, or object restored) to an SNS topic in the same Region. The SNS topic publishes the event to an SQS queue in the central Region.

The SQS queue is configured as the event source for your Lambda function and buffers the event messages for the Lambda function.

The Lambda function polls the SQS queue for messages and processes the Amazon S3 event notifications according to your application's requirements.

<https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/subscribe-a-lambda-function-to-event-notifications-from-s3-buckets-in-different-aws-regions.html>

NEW QUESTION 132

- (Topic 1)

A company is preparing to store confidential data in Amazon S3. For compliance reasons the data must be encrypted at rest. Encryption key usage must be logged for auditing purposes. Keys must be rotated every year.

Which solution meets these requirements and is the MOST operationally efficient?

- A. Server-side encryption with customer-provided keys (SSE-C)
- B. Server-side encryption with Amazon S3 managed keys (SSE-S3)
- C. Server-side encryption with AWS KMS (SSE-KMS) customer master keys (CMKs) with manual rotation
- D. Server-side encryption with AWS KMS (SSE-KMS) customer master keys (CMKs) with automatic rotation

Answer: D

Explanation:

<https://docs.aws.amazon.com/kms/latest/developerguide/rotate-keys.html> When you enable automatic key rotation for a customer managed key, AWS KMS generates new cryptographic material for the KMS key every year. AWS KMS also saves the KMS key's older cryptographic material in perpetuity so it can be used to decrypt data that the KMS key encrypted.

Key rotation in AWS KMS is a cryptographic best practice that is designed to be transparent and easy to use. AWS KMS supports optional automatic key rotation only for customer managed CMKs. Enable and disable key rotation. Automatic key rotation is disabled by default on customer managed CMKs. When you enable (or re-enable) key rotation, AWS KMS automatically rotates the CMK 365 days after the enable date and every 365 days thereafter.

NEW QUESTION 134

- (Topic 1)

A solutions architect is designing a VPC with public and private subnets. The VPC and subnets use IPv4 CIDR blocks. There is one public subnet and one private subnet in each of three Availability Zones (AZs) for high availability. An internet gateway is used to provide internet access for the public subnets. The private subnets require access to the internet to allow Amazon EC2 instances to download software updates. What should the solutions architect do to enable Internet access for the private subnets?

- A. Create three NAT gateways, one for each public subnet in each A
- B. Create a private route table for each AZ that forwards non-VPC traffic to the NAT gateway in its AZ.
- C. Create three NAT instances, one for each private subnet in each A
- D. Create a private route table for each AZ that forwards non-VPC traffic to the NAT instance in its AZ.
- E. Create a second internet gateway on one of the private subnet
- F. Update the route table for the private subnets that forward non-VPC traffic to the private internet gateway.
- G. Create an egress-only internet gateway on one of the public subnet
- H. Update the route table for the private subnets that forward non-VPC traffic to the egress- only internet gateway.

Answer: A

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2018/03/introducing-amazon-vpc-nat-gateway-in-the-aws-govcloud-us-region/#:~:text=NAT%20Gateway%20is%20a%20highly,instances%20in%20a%20private%20subnet.>
<https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-comparison.html>

NEW QUESTION 135

- (Topic 1)

A solutions architect is designing a two-tier web application. The application consists of a public-facing web tier hosted on Amazon EC2 in public subnets. The database tier consists of Microsoft SQL Server running on Amazon EC2 in a private subnet. Security is a high priority for the company. How should security groups be configured in this situation? (Select TWO)

- A. Configure the security group for the web tier to allow inbound traffic on port 443 from 0.0.0.0/0.
- B. Configure the security group for the web tier to allow outbound traffic on port 443 from 0.0.0.0/0.
- C. Configure the security group for the database tier to allow inbound traffic on port 1433 from the security group for the web tier.
- D. Configure the security group for the database tier to allow outbound traffic on ports 443 and 1433 to the security group for the web tier.
- E. Configure the security group for the database tier to allow inbound traffic on ports 443 and 1433 from the security group for the web tier.

Answer: AC

Explanation:

"Security groups create an outbound rule for every inbound rule." Not completely right. Stateful does NOT mean that if you create an inbound (or outbound) rule, it will create an outbound (or inbound) rule. What it does mean is: suppose you create an inbound rule on port 443 for the X ip. When a request enters on port 443 from X ip, it will allow traffic out for that request in the port 443. However, if you look at the outbound rules, there will not be any outbound rule on port 443 unless explicitly create it. In ACLs, which are stateless, you would have to create an inbound rule to allow incoming requests and an outbound rule to allow your application responds to those incoming requests.

https://docs.aws.amazon.com/vpc/latest/userguide/VPC_SecurityGroups.html#SecurityGroupRules

NEW QUESTION 136

- (Topic 1)

A company has an automobile sales website that stores its listings in a database on Amazon RDS. When an automobile is sold, the listing needs to be removed from the website and the data must be sent to multiple target systems. Which design should a solutions architect recommend?

- A. Create an AWS Lambda function triggered when the database on Amazon RDS is updated to send the information to an Amazon Simple Queue Service (Amazon SQS) queue for the targets to consume
- B. Create an AWS Lambda function triggered when the database on Amazon RDS is updated to send the information to an Amazon Simple Queue Service (Amazon SQS) FIFO queue for the targets to consume
- C. Subscribe to an RDS event notification and send an Amazon Simple Queue Service (Amazon SQS) queue fanned out to multiple Amazon Simple Notification Service (Amazon SNS) topics. Use AWS Lambda functions to update the targets
- D. Subscribe to an RDS event notification and send an Amazon Simple Notification Service (Amazon SNS) topic fanned out to multiple Amazon Simple Queue Service (Amazon SQS) queues. Use AWS Lambda functions to update the targets

Answer: D

Explanation:

<https://docs.aws.amazon.com/lambda/latest/dg/services-rds.html> <https://docs.aws.amazon.com/lambda/latest/dg/with-sns.html>

NEW QUESTION 141

- (Topic 1)

An application runs on an Amazon EC2 instance in a VPC. The application processes logs that are stored in an Amazon S3 bucket. The EC2 instance needs to access the S3 bucket without connectivity to the internet.

Which solution will provide private network connectivity to Amazon S3?

- A. Create a gateway VPC endpoint to the S3 bucket.
- B. Stream the logs to Amazon CloudWatch Log
- C. Export the logs to the S3 bucket.
- D. Create an instance profile on Amazon EC2 to allow S3 access.
- E. Create an Amazon API Gateway API with a private link to access the S3 endpoint.

Answer: A

Explanation:

VPC endpoint allows you to connect to AWS services using a private network instead of using the public Internet

NEW QUESTION 143

- (Topic 1)

A company has thousands of edge devices that collectively generate 1 TB of status alerts each day. Each alert is approximately 2 KB in size. A solutions architect needs to implement a solution to ingest and store the alerts for future analysis.

The company wants a highly available solution. However, the company needs to minimize costs and does not want to manage additional infrastructure. Additionally, the company wants to keep 14 days of data available for immediate analysis and archive any data older than 14 days.

What is the MOST operationally efficient solution that meets these requirements?

- A. Create an Amazon Kinesis Data Firehose delivery stream to ingest the alerts Configure the Kinesis Data Firehose stream to deliver the alerts to an Amazon S3 bucket Set up an S3 Lifecycle configuration to transition data to Amazon S3 Glacier after 14 days
- B. Launch Amazon EC2 instances across two Availability Zones and place them behind an Elastic Load Balancer to ingest the alerts Create a script on the EC2 instances that will store the alerts in an Amazon S3 bucket Set up an S3 Lifecycle configuration to transition data to Amazon S3 Glacier after 14 days
- C. Create an Amazon Kinesis Data Firehose delivery stream to ingest the alerts Configure the Kinesis Data Firehose stream to deliver the alerts to an Amazon Elasticsearch Service (Amazon ES) cluster Set up the Amazon ES cluster to take manual snapshots every day and delete data from the cluster that is older than 14 days
- D. Create an Amazon Simple Queue Service (Amazon SQS) standard queue to ingest the alerts and set the message retention period to 14 days Configure consumers to poll the SQS queue check the age of the message and analyze the message data as needed If the message is 14 days old the consumer should copy the message to an Amazon S3 bucket and delete the message from the SQS queue

Answer: A

Explanation:

<https://aws.amazon.com/kinesis/data-firehose/features/?nc=sn&loc=2#:~:text=into%20Amazon%20S3%2C%20Amazon%20Redshift%2C%20Amazon%20OpenSearch%20Service%2C%20Kinesis,Delivery%20streams>

NEW QUESTION 147

- (Topic 1)

A company has a production web application in which users upload documents through a web interface or a mobile app. According to a new regulatory requirement, new documents cannot be modified or deleted after they are stored.

What should a solutions architect do to meet this requirement?

- A. Store the uploaded documents in an Amazon S3 bucket with S3 Versioning and S3 Object Lock enabled
- B. Store the uploaded documents in an Amazon S3 bucket
- C. Configure an S3 Lifecycle policy to archive the documents periodically.
- D. Store the uploaded documents in an Amazon S3 bucket with S3 Versioning enabled Configure an ACL to restrict all access to read-only.
- E. Store the uploaded documents on an Amazon Elastic File System (Amazon EFS) volume
- F. Access the data by mounting the volume in read-only mode.

Answer: A

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock-overview.html>

NEW QUESTION 148

- (Topic 1)

A company has an application that runs on Amazon EC2 instances and uses an Amazon

Aurora database. The EC2 instances connect to the database by using user names and passwords that are stored locally in a file. The company wants to minimize the operational overhead of credential management.

What should a solutions architect do to accomplish this goal?

- A. Use AWS Secrets Manager
- B. Turn on automatic rotation.
- C. Use AWS Systems Manager Parameter Store
- D. Turn on automatic rotation.
- E. Create an Amazon S3 bucket to store objects that are encrypted with an AWS Key Management Service (AWS KMS) encryption key
- F. Migrate the credential file to the S3 bucket
- G. Point the application to the S3 bucket.
- H. Create an encrypted Amazon Elastic Block Store (Amazon EBS) volume (or each EC2 instance)
- I. Attach the new EBS volume to each EC2 instance
- J. Migrate the credential file to the new EBS volume
- K. Point the application to the new EBS volume.

Answer: A

Explanation:

<https://aws.amazon.com/cn/blogs/security/how-to-connect-to-aws-secrets-manager-service-within-a-virtual-private-cloud/>
<https://aws.amazon.com/blogs/security/rotate-amazon-rds-database-credentials-automatically-with-aws-secrets-manager/>

NEW QUESTION 151

- (Topic 1)

A company uses AWS Organizations to manage multiple AWS accounts for different departments. The management account has an Amazon S3 bucket that contains project reports. The company wants to limit access to this S3 bucket to only users of accounts within the organization in AWS Organizations.

Which solution meets these requirements with the LEAST amount of operational overhead?

- A. Add the aws:PrincipalOrgID global condition key with a reference to the organization ID to the S3 bucket policy.
- B. Create an organizational unit (OU) for each department
- C. Add the aws:PrincipalOrgPaths global condition key to the S3 bucket policy.
- D. Use AWS CloudTrail to monitor the CreateAccount, InviteAccountToOrganization, LeaveOrganization, and RemoveAccountFromOrganization event

- E. Update the S3 bucket policy accordingly.
- F. Tag each user that needs access to the S3 bucket.
- G. Add the aws:PrincipalTag global condition key to the S3 bucket policy.

Answer: A

Explanation:

<https://aws.amazon.com/blogs/security/control-access-to-aws-resources-by-using-the-aws-organization-of-iam-principals/>

The aws:PrincipalOrgID global key provides an alternative to listing all the account IDs for all AWS accounts in an organization. For example, the following Amazon S3 bucket policy allows members of any account in the XXX organization to add an object into the examtopics bucket.

```
{
  "Version": "2020-09-10",
  "Statement": [
    {
      "Sid": "AllowPutObject",
      "Effect": "Allow",
      "Principal": "*",
      "Action": "s3:PutObject",
      "Resource": "arn:aws:s3:::examtopics/*",
      "Condition": {
        "StringEquals": {
          "aws:PrincipalOrgID": ["XXX"]
        }
      }
    }
  ]
}
```

https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_condition-keys.html

NEW QUESTION 152

- (Topic 1)

A company runs its Infrastructure on AWS and has a registered base of 700,000 users for a document management application. The company intends to create a product that converts large PDF files to JPG image files. The PDF files average 5 MB in size. The company needs to store the original files and the converted files. A solutions architect must design a scalable solution to accommodate demand that will grow rapidly over time.

Which solution meets these requirements MOST cost-effectively?

- A. Save the PDF files to Amazon S3. Configure an S3 PUT event to invoke an AWS Lambda function to convert the files to JPG format and store them back in Amazon S3.
- B. Save the PDF files to Amazon DynamoDB.
- C. Use the DynamoDB Streams feature to invoke an AWS Lambda function to convert the files to JPG format and store them back in DynamoDB.
- D. Upload the PDF files to an AWS Elastic Beanstalk application that includes Amazon EC2 instances.
- E. Amazon Elastic Block Store (Amazon EBS) storage and an Auto Scaling group.
- F. Use a program in the EC2 instances to convert the files to JPG format. Save the PDF files and the JPG files in the EBS store.
- G. Upload the PDF files to an AWS Elastic Beanstalk application that includes Amazon EC2 instances, Amazon Elastic File System (Amazon EFS) storage, and an Auto Scaling group.
- H. Use a program in the EC2 instances to convert the files to JPG format. Save the PDF files and the JPG files in the EBS store.

Answer: A

Explanation:

Elastic Beanstalk is expensive, and DocumentDB has a 400KB max to upload files. So Lambda and S3 should be the one.

NEW QUESTION 153

- (Topic 1)

A company is running a business-critical web application on Amazon EC2 instances behind an Application Load Balancer. The EC2 instances are in an Auto Scaling group. The application uses an Amazon Aurora PostgreSQL database that is deployed in a single Availability Zone. The company wants the application to be highly available with minimum downtime and minimum loss of data.

Which solution will meet these requirements with the LEAST operational effort?

- A. Place the EC2 instances in different AWS Regions.
- B. Use Amazon Route 53 health checks to redirect traffic.
- C. Use Amazon Aurora PostgreSQL Cross-Region Replication.
- D. Configure the Auto Scaling group to use multiple Availability Zones.
- E. Configure the database as Multi-AZ.
- F. Configure an Amazon RDS Proxy instance for the database.
- G. Configure the Auto Scaling group to use one Availability Zone.
- H. Generate hourly snapshots of the database.
- I. Recover the database from the snapshots in the event of a failure.
- J. Configure the Auto Scaling group to use multiple AWS Regions.
- K. Write the data from the application to Amazon S3. Use S3 Event Notifications to launch an AWS Lambda function to write the data to the database.

Answer: B

Explanation:

To achieve high availability with minimum downtime and minimum loss of data, the Auto Scaling group should be configured to use multiple Availability Zones to ensure that there is no single point of failure. The database should be configured as Multi-AZ to enable automatic failover in case of an outage in the primary Availability Zone. Additionally, an Amazon RDS Proxy instance can be used to improve the scalability and availability of the database by reducing connection failures and improving failover times.

NEW QUESTION 158

- (Topic 2)

A company hosts a two-tier application on Amazon EC2 instances and Amazon RDS. The application's demand varies based on the time of day. The load is minimal after work hours and on weekends. The EC2 instances run in an EC2 Auto Scaling group that is configured with a minimum of two instances and a maximum of five instances. The application must be available at all times, but the company is concerned about overall cost.

Which solution meets the availability requirement MOST cost-effectively?

- A. Use all EC2 Spot Instances.
- B. Stop the RDS database when it is not in use.
- C. Purchase EC2 Instance Savings Plans to cover five EC2 instances.

- D. Purchase an RDS Reserved DB Instance
- E. Purchase two EC2 Reserved Instances Use up to three additional EC2 Spot Instances as needed
- F. Stop the RDS database when it is not in use.
- G. Purchase EC2 Instance Savings Plans to cover two EC2 instances
- H. Use up to three additional EC2 On-Demand Instances as needed
- I. Purchase an RDS Reserved DB Instance.

Answer: C

Explanation:

This solution meets the requirements of a two-tier application that has a variable demand based on the time of day and must be available at all times, while minimizing the overall cost. EC2 Reserved Instances can provide significant savings compared to On-Demand Instances for the baseline level of usage, and they can guarantee capacity reservation when needed. EC2 Spot Instances can provide up to 90% savings compared to On-Demand Instances for any additional capacity that the application needs during peak hours. Spot Instances are suitable for stateless applications that can tolerate interruptions and can be replaced by other instances. Stopping the RDS database when it is not in use can reduce the cost of running the database tier.

Option A is incorrect because using all EC2 Spot Instances can affect the availability of the application if there are not enough spare capacity or if the Spot price exceeds the maximum price. Stopping the RDS database when it is not in use can reduce the cost of running the database tier, but it can also affect the availability of the application. Option B is incorrect because purchasing EC2 Instance Savings Plans to cover five EC2 instances can lock in a fixed amount of compute usage per hour, which may not match the actual usage pattern of the application. Purchasing an RDS Reserved DB Instance can provide savings for the database tier, but it does not allow stopping the database when it is not in use. Option D is incorrect because purchasing EC2 Instance Savings Plans to cover two EC2 instances can lock in a fixed amount of compute usage per hour, which may not match the actual usage pattern of the application. Using up to three additional EC2 On-Demand Instances as needed can incur higher costs than using Spot Instances.

References:

? <https://aws.amazon.com/ec2/pricing/reserved-instances/>

? <https://aws.amazon.com/ec2/spot/>

? https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_StopInstance.html

NEW QUESTION 163

- (Topic 2)

A company has a multi-tier application that runs six front-end web servers in an Amazon EC2 Auto Scaling group in a single Availability Zone behind an Application Load Balancer (ALB). A solutions architect needs to modify the infrastructure to be highly available without modifying the application.

Which architecture should the solutions architect choose that provides high availability?

- A. Create an Auto Scaling group that uses three Instances across each of two Regions.
- B. Modify the Auto Scaling group to use three instances across each of two Availability Zones.
- C. Create an Auto Scaling template that can be used to quickly create more instances in another Region.
- D. Change the ALB in front of the Amazon EC2 instances in a round-robin configuration to balance traffic to the web tier.

Answer: B

Explanation:

High availability can be enabled for this architecture quite simply by modifying the existing Auto Scaling group to use multiple availability zones. The ASG will automatically balance the load so you don't actually need to specify the instances per AZ.

NEW QUESTION 164

- (Topic 2)

A company produces batch data that comes from different databases. The company also produces live stream data from network sensors and application APIs. The company needs to consolidate all the data into one place for business analytics. The company needs to process the incoming data and then stage the data in different Amazon S3 buckets. Teams will later run one-time queries and import the data into a business intelligence tool to show key performance indicators (KPIs).

Which combination of steps will meet these requirements with the LEAST operational overhead? (Choose two.)

- A. Use Amazon Athena for one-time queries Use Amazon QuickSight to create dashboards for KPIs
- B. Use Amazon Kinesis Data Analytics for one-time queries Use Amazon QuickSight to create dashboards for KPIs
- C. Create custom AWS Lambda functions to move the individual records from the databases to an Amazon Redshift cluster
- D. Use an AWS Glue extract transform, and load (ETL) job to convert the data into JSON format Load the data into multiple Amazon OpenSearch Service (Amazon Elasticsearch Service) clusters
- E. Use blueprints in AWS Lake Formation to identify the data that can be ingested into a data lake Use AWS Glue to crawl the source extract the data and load the data into Amazon S3 in Apache Parquet format

Answer: AE

Explanation:

Amazon Athena is the best choice for running one-time queries on streaming data. Although Amazon Kinesis Data Analytics provides an easy and familiar standard SQL language to analyze streaming data in real-time, it is designed for continuous queries rather than one-time queries[1]. On the other hand, Amazon Athena is a serverless interactive query service that allows querying data in Amazon S3 using SQL. It is optimized for ad-hoc querying and is ideal for running one-time queries on streaming data[2]. AWS Lake Formation uses as a central place to have all your data for analytics purposes (E). Athena integrates perfectly with S3 and can make queries (A).

NEW QUESTION 168

- (Topic 2)

A solutions architect needs to help a company optimize the cost of running an application on AWS. The application will use Amazon EC2 instances, AWS Fargate, and AWS Lambda for compute within the architecture.

The EC2 instances will run the data ingestion layer of the application. EC2 usage will be sporadic and unpredictable. Workloads that run on EC2 instances can be interrupted at any time. The application front end will run on Fargate, and Lambda will serve the API layer. The front-end utilization and API layer utilization will be predictable over the course of the next year.

Which combination of purchasing options will provide the MOST cost-effective solution for hosting this application? (Choose two.)

- A. Use Spot Instances for the data ingestion layer

- B. Use On-Demand Instances for the data ingestion layer
- C. Purchase a 1-year Compute Savings Plan for the front end and API layer.
- D. Purchase 1-year All Upfront Reserved instances for the data ingestion layer.
- E. Purchase a 1-year EC2 instance Savings Plan for the front end and API layer.

Answer: AC

Explanation:

EC2 instance Savings Plan saves 72% while Compute Savings Plans saves 66%. But according to link, it says "Compute Savings Plans provide the most flexibility and help to reduce your costs by up to 66%. These plans automatically apply to EC2 instance usage regardless of instance family, size, AZ, region, OS or tenancy, and also apply to Fargate and Lambda usage." EC2 instance Savings Plans are not applied to Fargate or Lambda

NEW QUESTION 169

- (Topic 2)

A company wants to use the AWS Cloud to make an existing application highly available and resilient. The current version of the application resides in the company's data center. The application recently experienced data loss after a database server crashed because of an unexpected power outage. The company needs a solution that avoids any single points of failure. The solution must give the application the ability to scale to meet user demand. Which solution will meet these requirements?

- A. Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zone
- B. Use an Amazon RDS DB instance in a Multi-AZ configuration.
- C. Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group in a single Availability Zone
- D. Deploy the database on an EC2 instance
- E. Enable EC2 Auto Recovery.
- F. Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zone
- G. Use an Amazon RDS DB instance with a read replica in a single Availability Zone
- H. Promote the read replica to replace the primary DB instance if the primary DB instance fails.
- I. Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones. Deploy the primary and secondary database servers on EC2 instances across multiple Availability Zones. Use Amazon Elastic Block Store (Amazon EBS) Multi-Attach to create shared storage between the instances.

Answer: A

Explanation:

Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones. Use an Amazon RDS DB instance in a Multi-AZ configuration. To make an existing application highly available and resilient while avoiding any single points of failure and giving the application the ability to scale to meet user demand, the best solution would be to deploy the application servers using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones and use an Amazon RDS DB instance in a Multi-AZ configuration. By using an Amazon RDS DB instance in a Multi-AZ configuration, the database is automatically replicated across multiple Availability Zones, ensuring that the database is highly available and can withstand the failure of a single Availability Zone. This provides fault tolerance and avoids any single points of failure.

NEW QUESTION 171

- (Topic 2)

An application runs on Amazon EC2 instances across multiple Availability Zones. The instances run in an Amazon EC2 Auto Scaling group behind an Application Load Balancer. The application performs best when the CPU utilization of the EC2 instances is at or near 40%. What should a solutions architect do to maintain the desired performance across all instances in the group?

- A. Use a simple scaling policy to dynamically scale the Auto Scaling group
- B. Use a target tracking policy to dynamically scale the Auto Scaling group
- C. Use an AWS Lambda function to update the desired Auto Scaling group capacity.
- D. Use scheduled scaling actions to scale up and scale down the Auto Scaling group

Answer: B

Explanation:

<https://docs.aws.amazon.com/autoscaling/application/userguide/application-auto-scaling-target-tracking.html>

NEW QUESTION 176

- (Topic 2)

A company uses a popular content management system (CMS) for its corporate website. However, the required patching and maintenance are burdensome. The company is redesigning its website and wants a new solution. The website will be updated four times a year and does not need to have any dynamic content available. The solution must provide high scalability and enhanced security. Which combination of changes will meet these requirements with the LEAST operational overhead? (Choose two.)

- A. Deploy an AWS WAF web ACL in front of the website to provide HTTPS functionality
- B. Create and deploy an AWS Lambda function to manage and serve the website content
- C. Create the new website and an Amazon S3 bucket. Deploy the website on the S3 bucket with static website hosting enabled
- D. Create the new website
- E. Deploy the website by using an Auto Scaling group of Amazon EC2 instances behind an Application Load Balancer.

Answer: AD

Explanation:

A -> We can configure CloudFront to require HTTPS from clients (enhanced security)
<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/using-https-viewers-to-cloudfront.html> D -> storing static website on S3 provides scalability and less operational overhead, then configuration of Application LB and EC2 instances (hence E is out)

NEW QUESTION 180

- (Topic 2)

A company is migrating its on-premises PostgreSQL database to Amazon Aurora PostgreSQL. The on-premises database must remain online and accessible during the migration. The Aurora database must remain synchronized with the on-premises database. Which combination of actions must a solutions architect take to meet these requirements? (Choose two.)

- A. Create an ongoing replication task.
- B. Create a database backup of the on-premises database
- C. Create an AWS Database Migration Service (AWS DMS) replication server
- D. Convert the database schema by using the AWS Schema Conversion Tool (AWS SCT).
- E. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to monitor the database synchronization

Answer: AC

Explanation:

AWS Database Migration Service supports homogeneous migrations such as Oracle to Oracle, as well as heterogeneous migrations between different database platforms, such as Oracle or Microsoft SQL Server to Amazon Aurora. With AWS Database Migration Service, you can also continuously replicate data with low latency from any supported source to any supported target. For example, you can replicate from multiple sources to Amazon Simple Storage Service (Amazon S3) to build a highly available and scalable data lake solution. You can also consolidate databases into a petabyte-scale data warehouse by streaming data to Amazon Redshift. Learn more about the supported source and target databases. <https://aws.amazon.com/dms/>

NEW QUESTION 183

- (Topic 2)

A company is building a web-based application running on Amazon EC2 instances in multiple Availability Zones. The web application will provide access to a repository of text documents totaling about 900 TB in size. The company anticipates that the web application will experience periods of high demand. A solutions architect must ensure that the storage component for the text documents can scale to meet the demand of the application at all times. The company is concerned about the overall cost of the solution.

Which storage solution meets these requirements MOST cost-effectively?

- A. Amazon Elastic Block Store (Amazon EBS)
- B. Amazon Elastic File System (Amazon EFS)
- C. Amazon Elasticsearch Service (Amazon ES)
- D. Amazon S3

Answer: D

Explanation:

Amazon S3 is cheapest and can be accessed from anywhere.

NEW QUESTION 188

- (Topic 2)

A company is concerned about the security of its public web application due to recent web attacks. The application uses an Application Load Balancer (ALB). A solutions architect must reduce the risk of DDoS attacks against the application.

What should the solutions architect do to meet this requirement?

- A. Add an Amazon Inspector agent to the ALB.
- B. Configure Amazon Macie to prevent attacks.
- C. Enable AWS Shield Advanced to prevent attacks.
- D. Configure Amazon GuardDuty to monitor the ALB.

Answer: C

Explanation:

AWS Shield Advanced provides expanded DDoS attack protection for your Amazon EC2 instances, Elastic Load Balancing load balancers, CloudFront distributions, Route 53 hosted zones, and AWS Global Accelerator standard accelerators. <https://docs.aws.amazon.com/waf/latest/developerguide/what-is-aws-waf.html>

NEW QUESTION 191

- (Topic 2)

A company wants to measure the effectiveness of its recent marketing campaigns. The company performs batch processing on csv files of sales data and stores the results in an Amazon S3 bucket once every hour. The S3 bucket contains billions of objects. The company runs one-time queries in Amazon Athena to determine which products are most popular on a particular date for a particular region. Queries sometimes fail or take longer than expected to finish.

Which actions should a solutions architect take to improve the query performance and reliability? (Select TWO.)

- A. Reduce the S3 object sizes to less than 126 MB
- B. Partition the data by date and region in Amazon S3
- C. Store the files as large, single objects in Amazon S3.
- D. Use Amazon Kinesis Data Analytics to run the Queries as part of the batch processing operation
- E. Use an AWS Glue extract, transform, and load (ETL) process to convert the csv files into Apache Parquet format.

Answer: BE

Explanation:

<https://aws.amazon.com/blogs/big-data/top-10-performance-tuning-tips-for-amazon-athena/>

This solution meets the requirements of measuring the effectiveness of marketing campaigns by performing batch processing on csv files of sales data and storing the results in an Amazon S3 bucket once every hour. An AWS Glue ETL process can use services such as AWS Glue or AWS Data Pipeline to extract data from S3, transform it into a more efficient format such as Apache Parquet, and load it back into S3. Apache Parquet is a columnar storage format that can improve the query performance and reliability of Athena by reducing the amount of data scanned, improving compression ratio, and enabling predicate pushdown.

NEW QUESTION 194

- (Topic 2)

A company has an ecommerce checkout workflow that writes an order to a database and calls a service to process the payment. Users are experiencing timeouts during the checkout process. When users resubmit the checkout form, multiple unique orders are created for the same desired transaction. How should a solutions architect refactor this workflow to prevent the creation of multiple orders?

- A. Configure the web application to send an order message to Amazon Kinesis Data Firehose
- B. Set the payment service to retrieve the message from Kinesis Data Firehose and process the order.
- C. Create a rule in AWS CloudTrail to invoke an AWS Lambda function based on the logged application path request Use Lambda to query the database, call the payment service, and pass in the order information.
- D. Store the order in the databas
- E. Send a message that includes the order number to Amazon Simple Notification Service (Amazon SNS). Set the payment service to pollAmazon SN
- F. retrieve the message, and process the order.
- G. Store the order in the databas
- H. Send a message that includes the order number to an Amazon Simple Queue Service (Amazon SQS) FIFO queu
- I. Set the payment service to retrieve the message and process the orde
- J. Delete the message from the queue.

Answer: D

Explanation:

This approach ensures that the order creation and payment processing steps are separate and atomic. By sending the order information to an SQS FIFO queue, the payment service can process the order one at a time and in the order they were received. If the payment service is unable to process an order, it can be retried later, preventing the creation of multiple orders. The deletion of the message from the queue after it is processed will prevent the same message from being processed multiple times.

NEW QUESTION 199

- (Topic 2)

A new employee has joined a company as a deployment engineer. The deployment engineer will be using AWS CloudFormation templates to create multiple AWS resources. A solutions architect wants the deployment engineer to perform job activities while following the principle of least privilege. Which steps should the solutions architect do in conjunction to reach this goal? (Select two.)

- A. Have the deployment engineer use AWS account root user credentials for performing AWS CloudFormation stack operations.
- B. Create a new IAM user for the deployment engineer and add the IAM user to a group that has the PowerUsers IAM policy attached.
- C. Create a new IAM user for the deployment engineer and add the IAM user to a group that has the Administrate/Access IAM policy attached.
- D. Create a new IAM User for the deployment engineer and add the IAM user to a group that has an IAM policy that allows AWS CloudFormation actions only.
- E. Create an IAM role for the deployment engineer to explicitly define the permissions specific to the AWS CloudFormation stack and launch stacks using Dial IAM role.

Answer: DE

Explanation:

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles.html https://docs.aws.amazon.com/IAM/latest/UserGuide/id_users.html

NEW QUESTION 204

- (Topic 2)

A company is running an online transaction processing (OLTP) workload on AWS. This workload uses an unencrypted Amazon RDS DB instance in a Multi-AZ deployment. Daily database snapshots are taken from this instance.

What should a solutions architect do to ensure the database and snapshots are always encrypted moving forward?

- A. Encrypt a copy of the latest DB snapsho
- B. Replace existing DB instance by restoring the encrypted snapshot
- C. Create a new encrypted Amazon Elastic Block Store (Amazon EBS) volume and copy the snapshots to it Enable encryption on the DB instance
- D. Copy the snapshots and enable encryption using AWS Key Management Service (AWS KMS) Restore encrypted snapshot to an existing DB instance
- E. Copy the snapshots to an Amazon S3 bucket that is encrypted using server-side encryption with AWS Key Management Service (AWS KMS) managed keys (SSE-KMS)

Answer: A

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_RestoreFromSnapshot.html#USER_RestoreFromSnapshot.CON
Under "Encrypt unencrypted resources" - <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSEncryption.html>

NEW QUESTION 207

- (Topic 2)

A company wants to move its application to a serverless solution. The serverless solution needs to analyze existing and new data by using SL. The company stores the data in an Amazon S3 bucket. The data requires encryption and must be replicated to a different AWS Region.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Create a new S3 bucke
- B. Load the data into the new S3 bucke
- C. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Regio
- D. Use server-side encryption with AWS KMS multi-Region kays (SSE-KMS). Use Amazon Athena to query the data.
- E. Create a new S3 bucke
- F. Load the data into the new S3 bucke
- G. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Regio
- H. Use server-side encryption with AWS KMS multi-Region keys (SSE-KMS). Use Amazon RDS to query the data.
- I. Load the data into the existing S3 bucke
- J. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Regio
- K. Use server-side encryption with Amazon S3managed encryption keys (SSE-S3). Use Amazon Athena to query the data.
- L. Load the data into the existing S3 bucke
- M. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Regio

N. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3). Use Amazon RDS to query the data.

Answer: A

Explanation:

This solution meets the requirements of a serverless solution, encryption, replication, and SQL analysis with the least operational overhead. Amazon Athena is a serverless interactive query service that can analyze data in S3 using standard SQL. S3 Cross-Region Replication (CRR) can replicate encrypted objects to an S3 bucket in another Region automatically. Server-side encryption with AWS KMS multi-Region keys (SSE-KMS) can encrypt the data at rest using keys that are replicated across multiple Regions. Creating a new S3 bucket can avoid potential conflicts with existing data or configurations.

Option B is incorrect because Amazon RDS is not a serverless solution and it cannot query data in S3 directly. Option C is incorrect because server-side encryption with Amazon S3 managed encryption keys (SSE-S3) does not use KMS keys and it does not support multi-Region replication. Option D is incorrect because Amazon RDS is not a serverless solution and it cannot query data in S3 directly. It is also incorrect for the same reason as option C. References:

? <https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication-walkthrough-4.html>

? <https://aws.amazon.com/blogs/storage/considering-four-different-replication-options-for-data-in-amazon-s3/>

? <https://docs.aws.amazon.com/AmazonS3/latest/userguide/UsingEncryption.html>

? <https://aws.amazon.com/athena/>

NEW QUESTION 210

- (Topic 2)

A solutions architect is creating a new Amazon CloudFront distribution for an application. Some of the information submitted by users is sensitive. The application uses HTTPS but needs another layer of security. The sensitive information should be protected throughout the entire application stack, and access to the information should be restricted to certain applications.

Which action should the solutions architect take?

- A. Configure a CloudFront signed URL.
- B. Configure a CloudFront signed cookie.
- C. Configure a CloudFront field-level encryption profile.
- D. Configure CloudFront and set the Origin Protocol Policy setting to HTTPS Only for the Viewer Protocol Policy.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/field-level-encryption.html>

"With Amazon CloudFront, you can enforce secure end-to-end connections to origin servers by using HTTPS. Field-level encryption adds an additional layer of security that lets you protect specific data throughout system processing so that only certain applications can see it."

NEW QUESTION 213

- (Topic 2)

A company runs workloads on AWS. The company needs to connect to a service from an external provider. The service is hosted in the provider's VPC. According to the company's security team, the connectivity must be private and must be restricted to the target service. The connection must be initiated only from the company's VPC.

Which solution will meet these requirements?

- A. Create a VPC peering connection between the company's VPC and the provider's VPC
- B. Update the route table to connect to the target service.
- C. Ask the provider to create a virtual private gateway in its VPC
- D. Use AWS PrivateLink to connect to the target service.
- E. Create a NAT gateway in a public subnet of the company's VPC
- F. Update the route table to connect to the target service.
- G. Ask the provider to create a VPC endpoint for the target service
- H. Use AWS PrivateLink to connect to the target service.

Answer: D

Explanation:

AWS PrivateLink provides private connectivity between VPCs, AWS services, and your on-premises networks, without exposing your traffic to the public internet. AWS PrivateLink makes it easy to connect services across different accounts and VPCs to significantly simplify your network architecture. Interface **VPC endpoints**, powered by AWS PrivateLink, connect you to services hosted by AWS Partners and supported solutions available in AWS Marketplace.

<https://aws.amazon.com/privatelink/>

NEW QUESTION 216

- (Topic 2)

A company has an AWS account used for software engineering. The AWS account has access to the company's on-premises data center through a pair of AWS Direct Connect connections. All non-VPC traffic routes to the virtual private gateway.

A development team recently created an AWS Lambda function through the console. The development team needs to allow the function to access a database that runs in a private subnet in the company's data center.

Which solution will meet these requirements?

- A. Configure the Lambda function to run in the VPC with the appropriate security group.
- B. Set up a VPN connection from AWS to the data center
- C. Route the traffic from the Lambda function through the VPN.
- D. Update the route tables in the VPC to allow the Lambda function to access the on-premises data center through Direct Connect.
- E. Create an Elastic IP address
- F. Configure the Lambda function to send traffic through the Elastic IP address without an elastic network interface.

Answer: A

Explanation:

<https://docs.aws.amazon.com/lambda/latest/dg/configuration-vpc.html#vpc-managing-eni>

NEW QUESTION 221

- (Topic 2)

A company runs a production application on a fleet of Amazon EC2 instances. The application reads the data from an Amazon SQS queue and processes the messages in parallel. The message volume is unpredictable and often has intermittent traffic. This application should continually process messages without any downtime.

Which solution meets these requirements MOST cost-effectively?

- A. Use Spot Instances exclusively to handle the maximum capacity required.
- B. Use Reserved Instances exclusively to handle the maximum capacity required.
- C. Use Reserved Instances for the baseline capacity and use Spot Instances to handle additional capacity.
- D. Use Reserved Instances for the baseline capacity and use On-Demand Instances to handle additional capacity.

Answer: D

Explanation:

We recommend that you use On-Demand Instances for applications with short-term, irregular workloads that cannot be interrupted.
<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-on-demand-instances.html>

NEW QUESTION 225

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your SAA-C03 Exam with Our Prep Materials Via below:

<https://www.certleader.com/SAA-C03-dumps.html>